



(12) **United States Patent**  
**Hillman**

(10) **Patent No.:** **US 11,417,019 B2**  
(45) **Date of Patent:** **\*Aug. 16, 2022**

(54) **DETECTION OF A CALIBRATION OBJECT FOR MODIFYING IMAGE PARAMETERS**

(71) Applicant: **Unity Technologies SF**, San Francisco, CA (US)

(72) Inventor: **Peter M. Hillman**, Wellington (NZ)

(73) Assignee: **UNITY TECHNOLOGIES SF**, San Francisco, CA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/490,421**

(22) Filed: **Sep. 30, 2021**

(65) **Prior Publication Data**

US 2022/0020176 A1 Jan. 20, 2022

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 17/065,439, filed on Oct. 7, 2020.

(Continued)

(51) **Int. Cl.**

**G06T 7/80** (2017.01)

**G06T 7/90** (2017.01)

(Continued)

(52) **U.S. Cl.**

CPC ..... **G06T 7/80** (2017.01); **G06T 7/11** (2017.01); **G06T 7/162** (2017.01); **G06T 7/90** (2017.01); **G06T 2207/10024** (2013.01)

(58) **Field of Classification Search**

CPC ... **G06T 7/80**; **G06T 7/11**; **G06T 7/162**; **G06T 7/90**; **G06T 2207/10024**

(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,476,394 B1 \* 11/2002 Amitani ..... G01T 1/2018 250/368

2013/0336583 A1 \* 12/2013 Ernst ..... G06T 7/80 382/165

(Continued)

FOREIGN PATENT DOCUMENTS

EP 2 940 626 A2 11/2015

OTHER PUBLICATIONS

<https://xritephoto.com>—Colorchecker Classic—Camera & Image Calibration: X-Rate Photo & Video—4 pages.


*Primary Examiner* — Ali Bayat

(74) *Attorney, Agent, or Firm* — Trellis IP Law Group, PC

(57) **ABSTRACT**

Implementations provide for automated detection of a calibration object within a recorded image. In some implementations, a system receives an original image from a camera, wherein the original image includes at least a portion of a calibration chart. The system further derives a working image from the original image. The system further determines regions in the working image, wherein the regions include groups of pixels having values within a predetermined criterion. The system further analyzes two or more of the regions to identify a candidate calibration chart in the working image. The system further identifies at least one region within the candidate calibration chart as a patch, where the identifying of the at least one region is based on a color of the patch. The system further predicts a location of one or more additional patches based on at least the identified patch.

**20 Claims, 10 Drawing Sheets**



**Related U.S. Application Data**

(60) Provisional application No. 63/047,868, filed on Jul. 2, 2020.

(51) **Int. Cl.**

**G06T 7/162** (2017.01)

**G06T 7/11** (2017.01)

(58) **Field of Classification Search**

USPC ..... 382/162

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

2014/0211022 A1\* 7/2014 Koh ..... H04N 17/02  
348/188  
2014/0286569 A1\* 9/2014 Devic ..... G06T 7/73  
382/165  
2018/0080829 A1\* 3/2018 Tanimura ..... G01N 21/251  
2019/0066338 A1 2/2019 Perlman et al.  
2019/0306477 A1\* 10/2019 Nordback ..... G06T 7/90

\* cited by examiner




FIG. 1

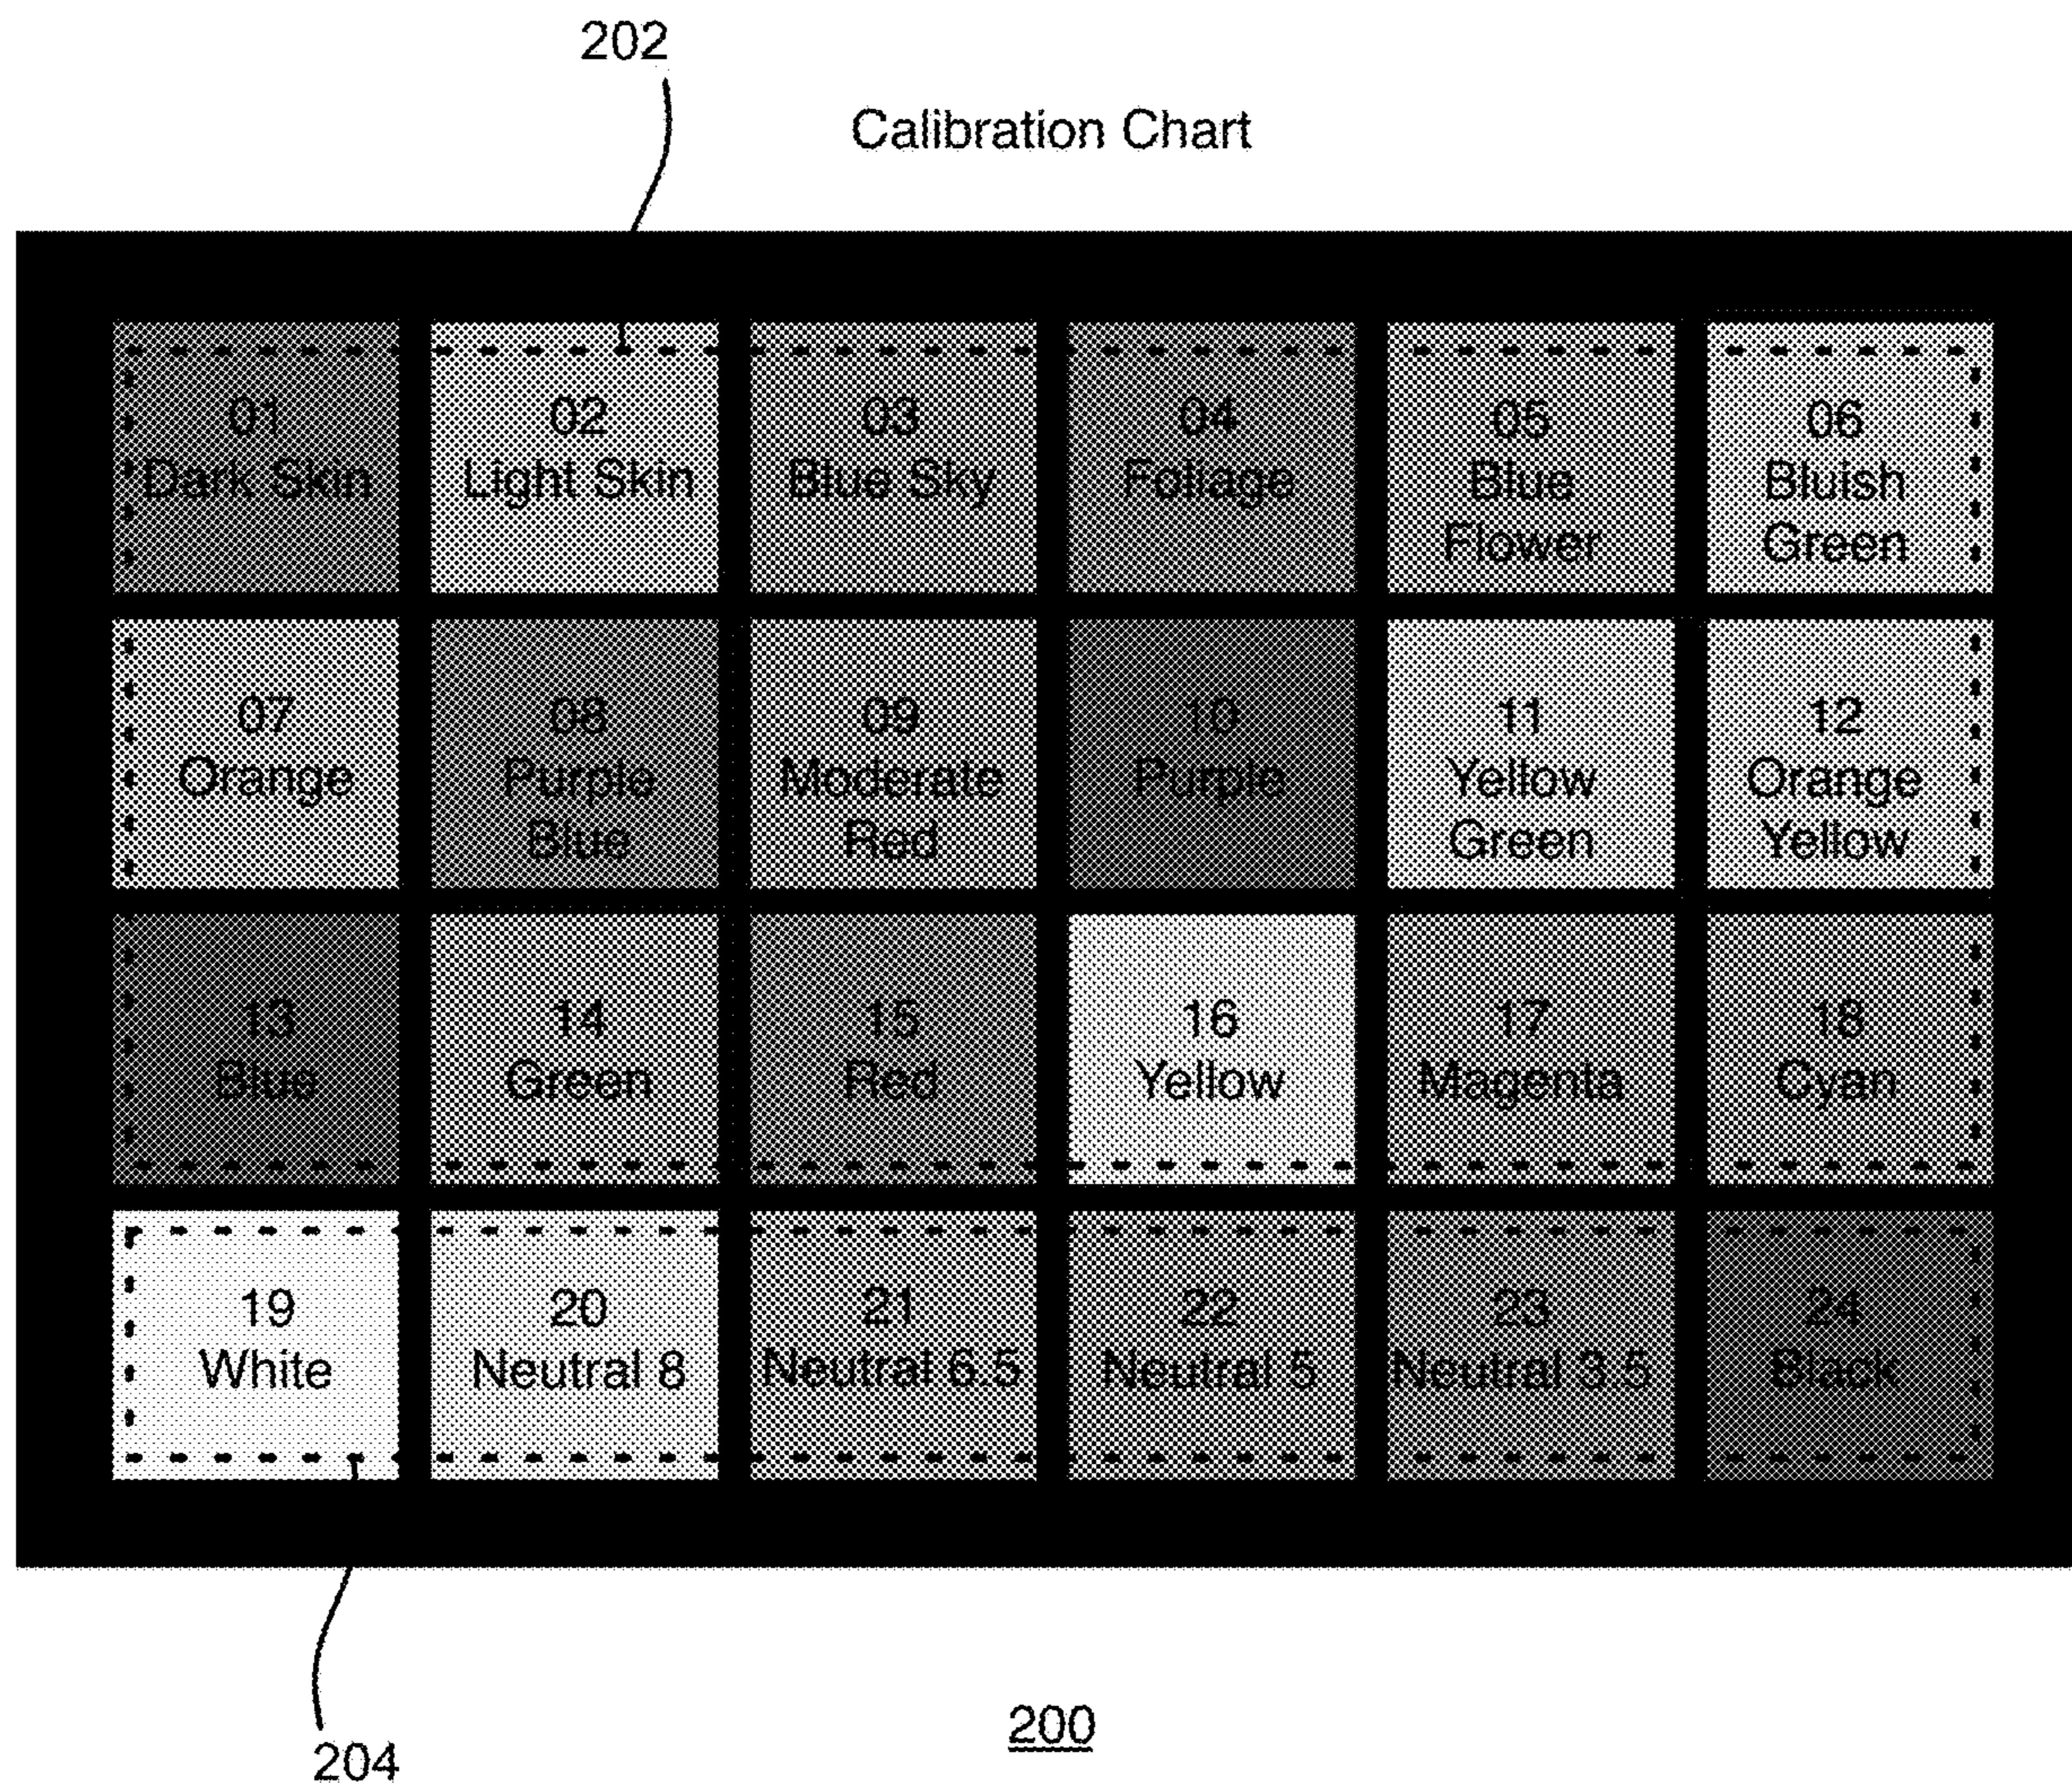


FIG. 2

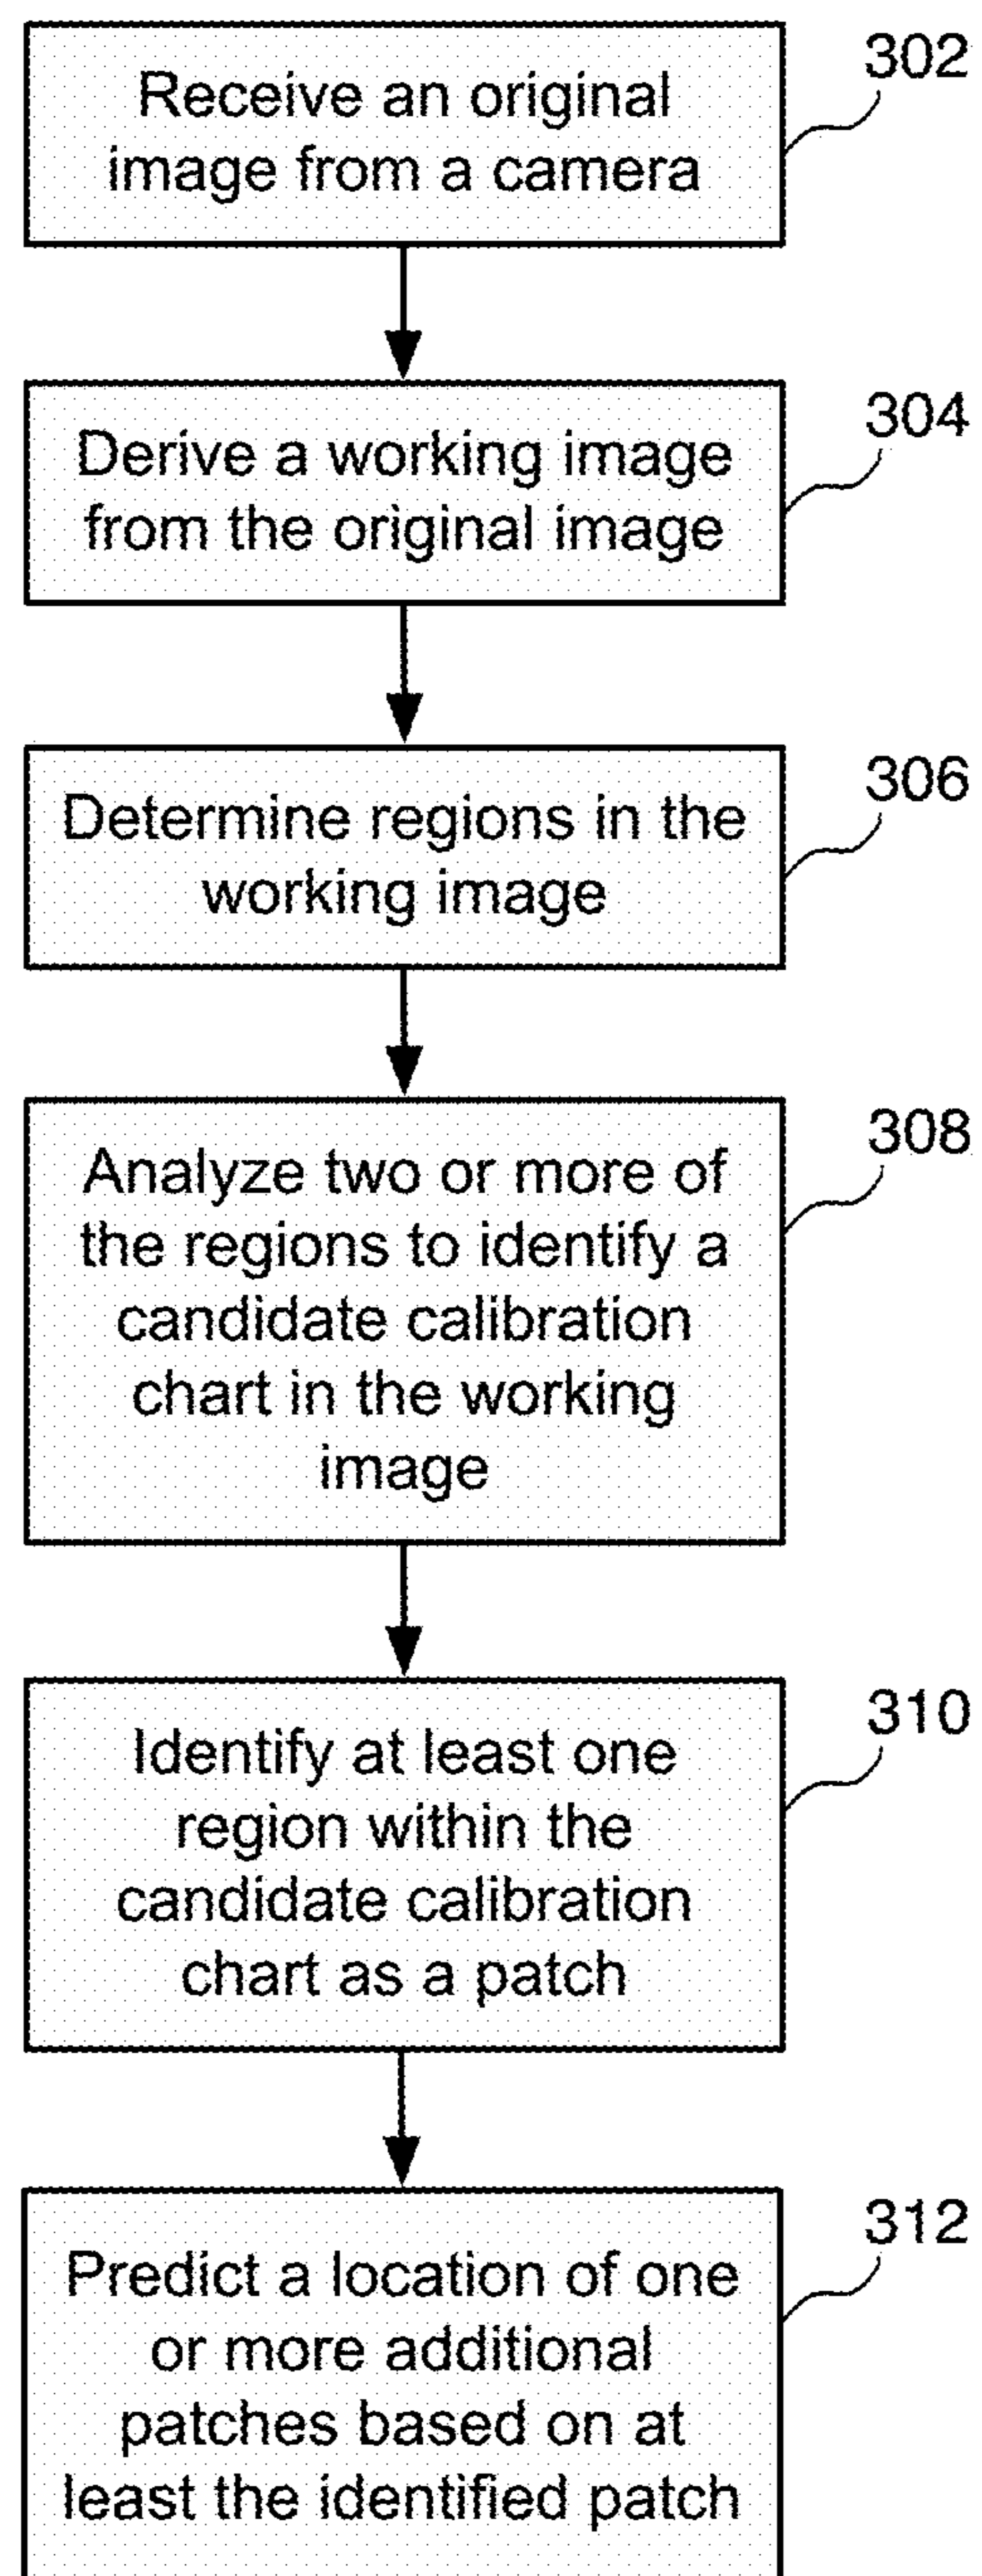




FIG. 3




400

FIG. 4



500  
FIG. 5



600

FIG. 6



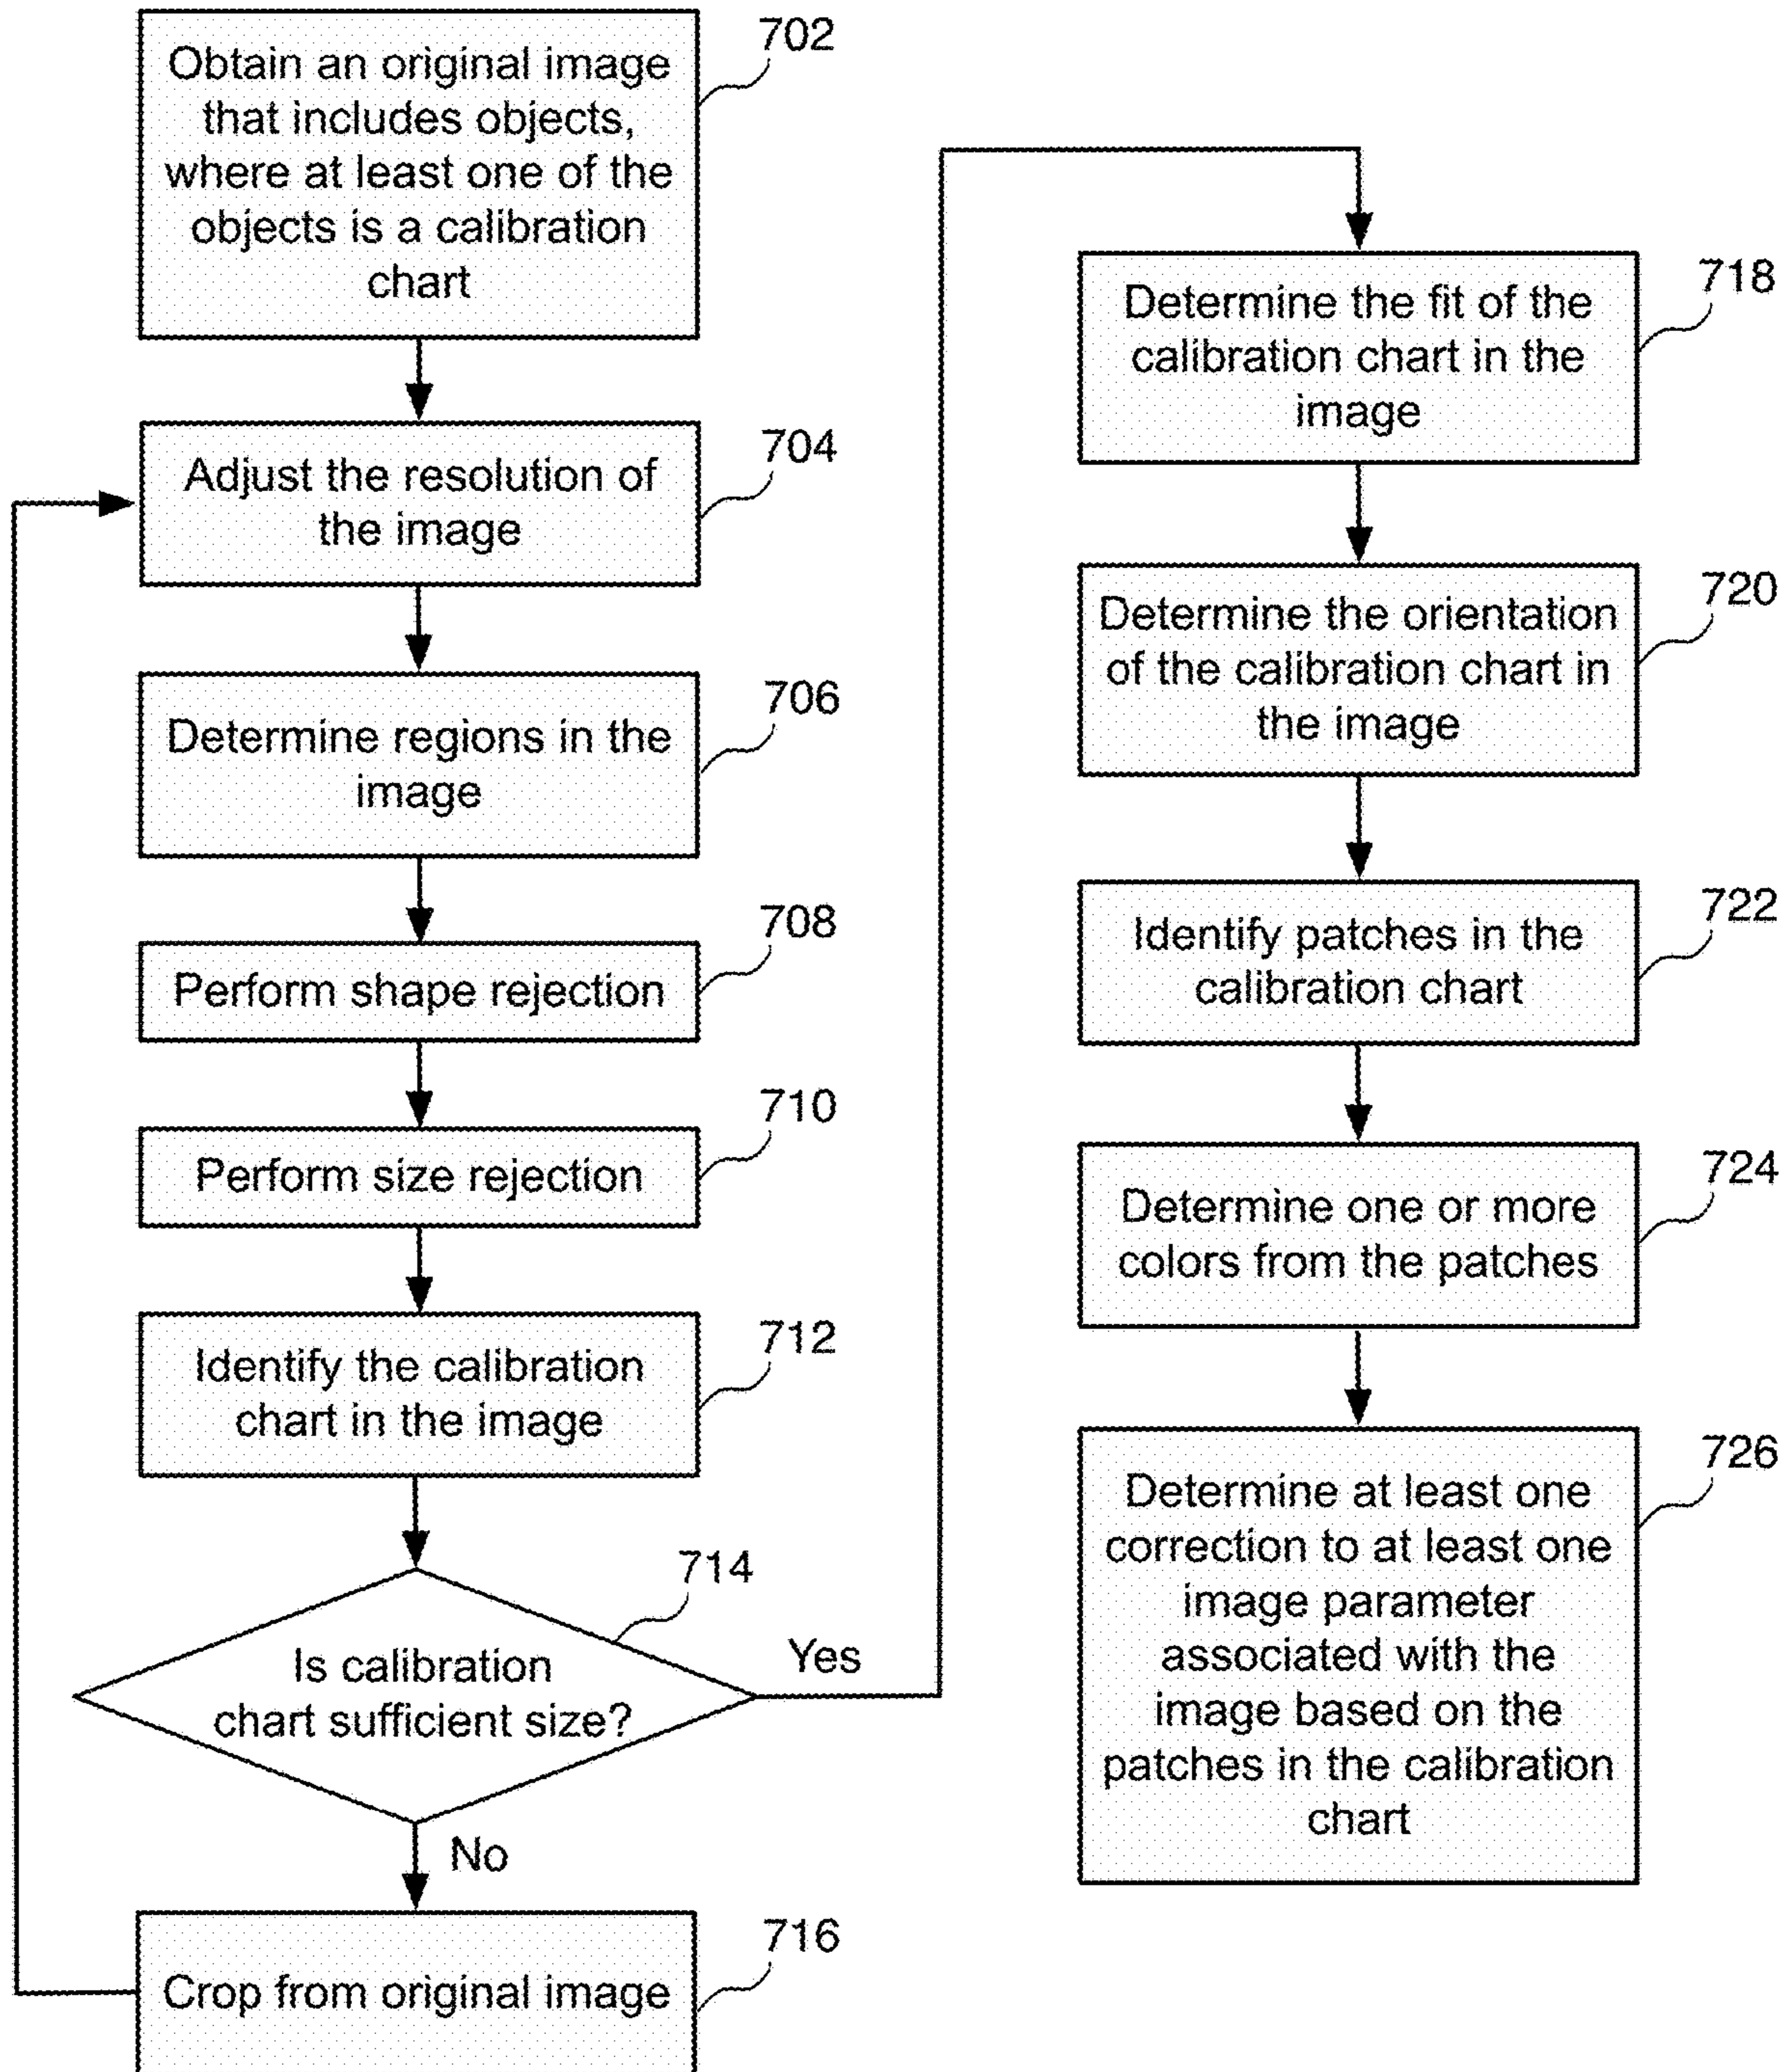


FIG. 7

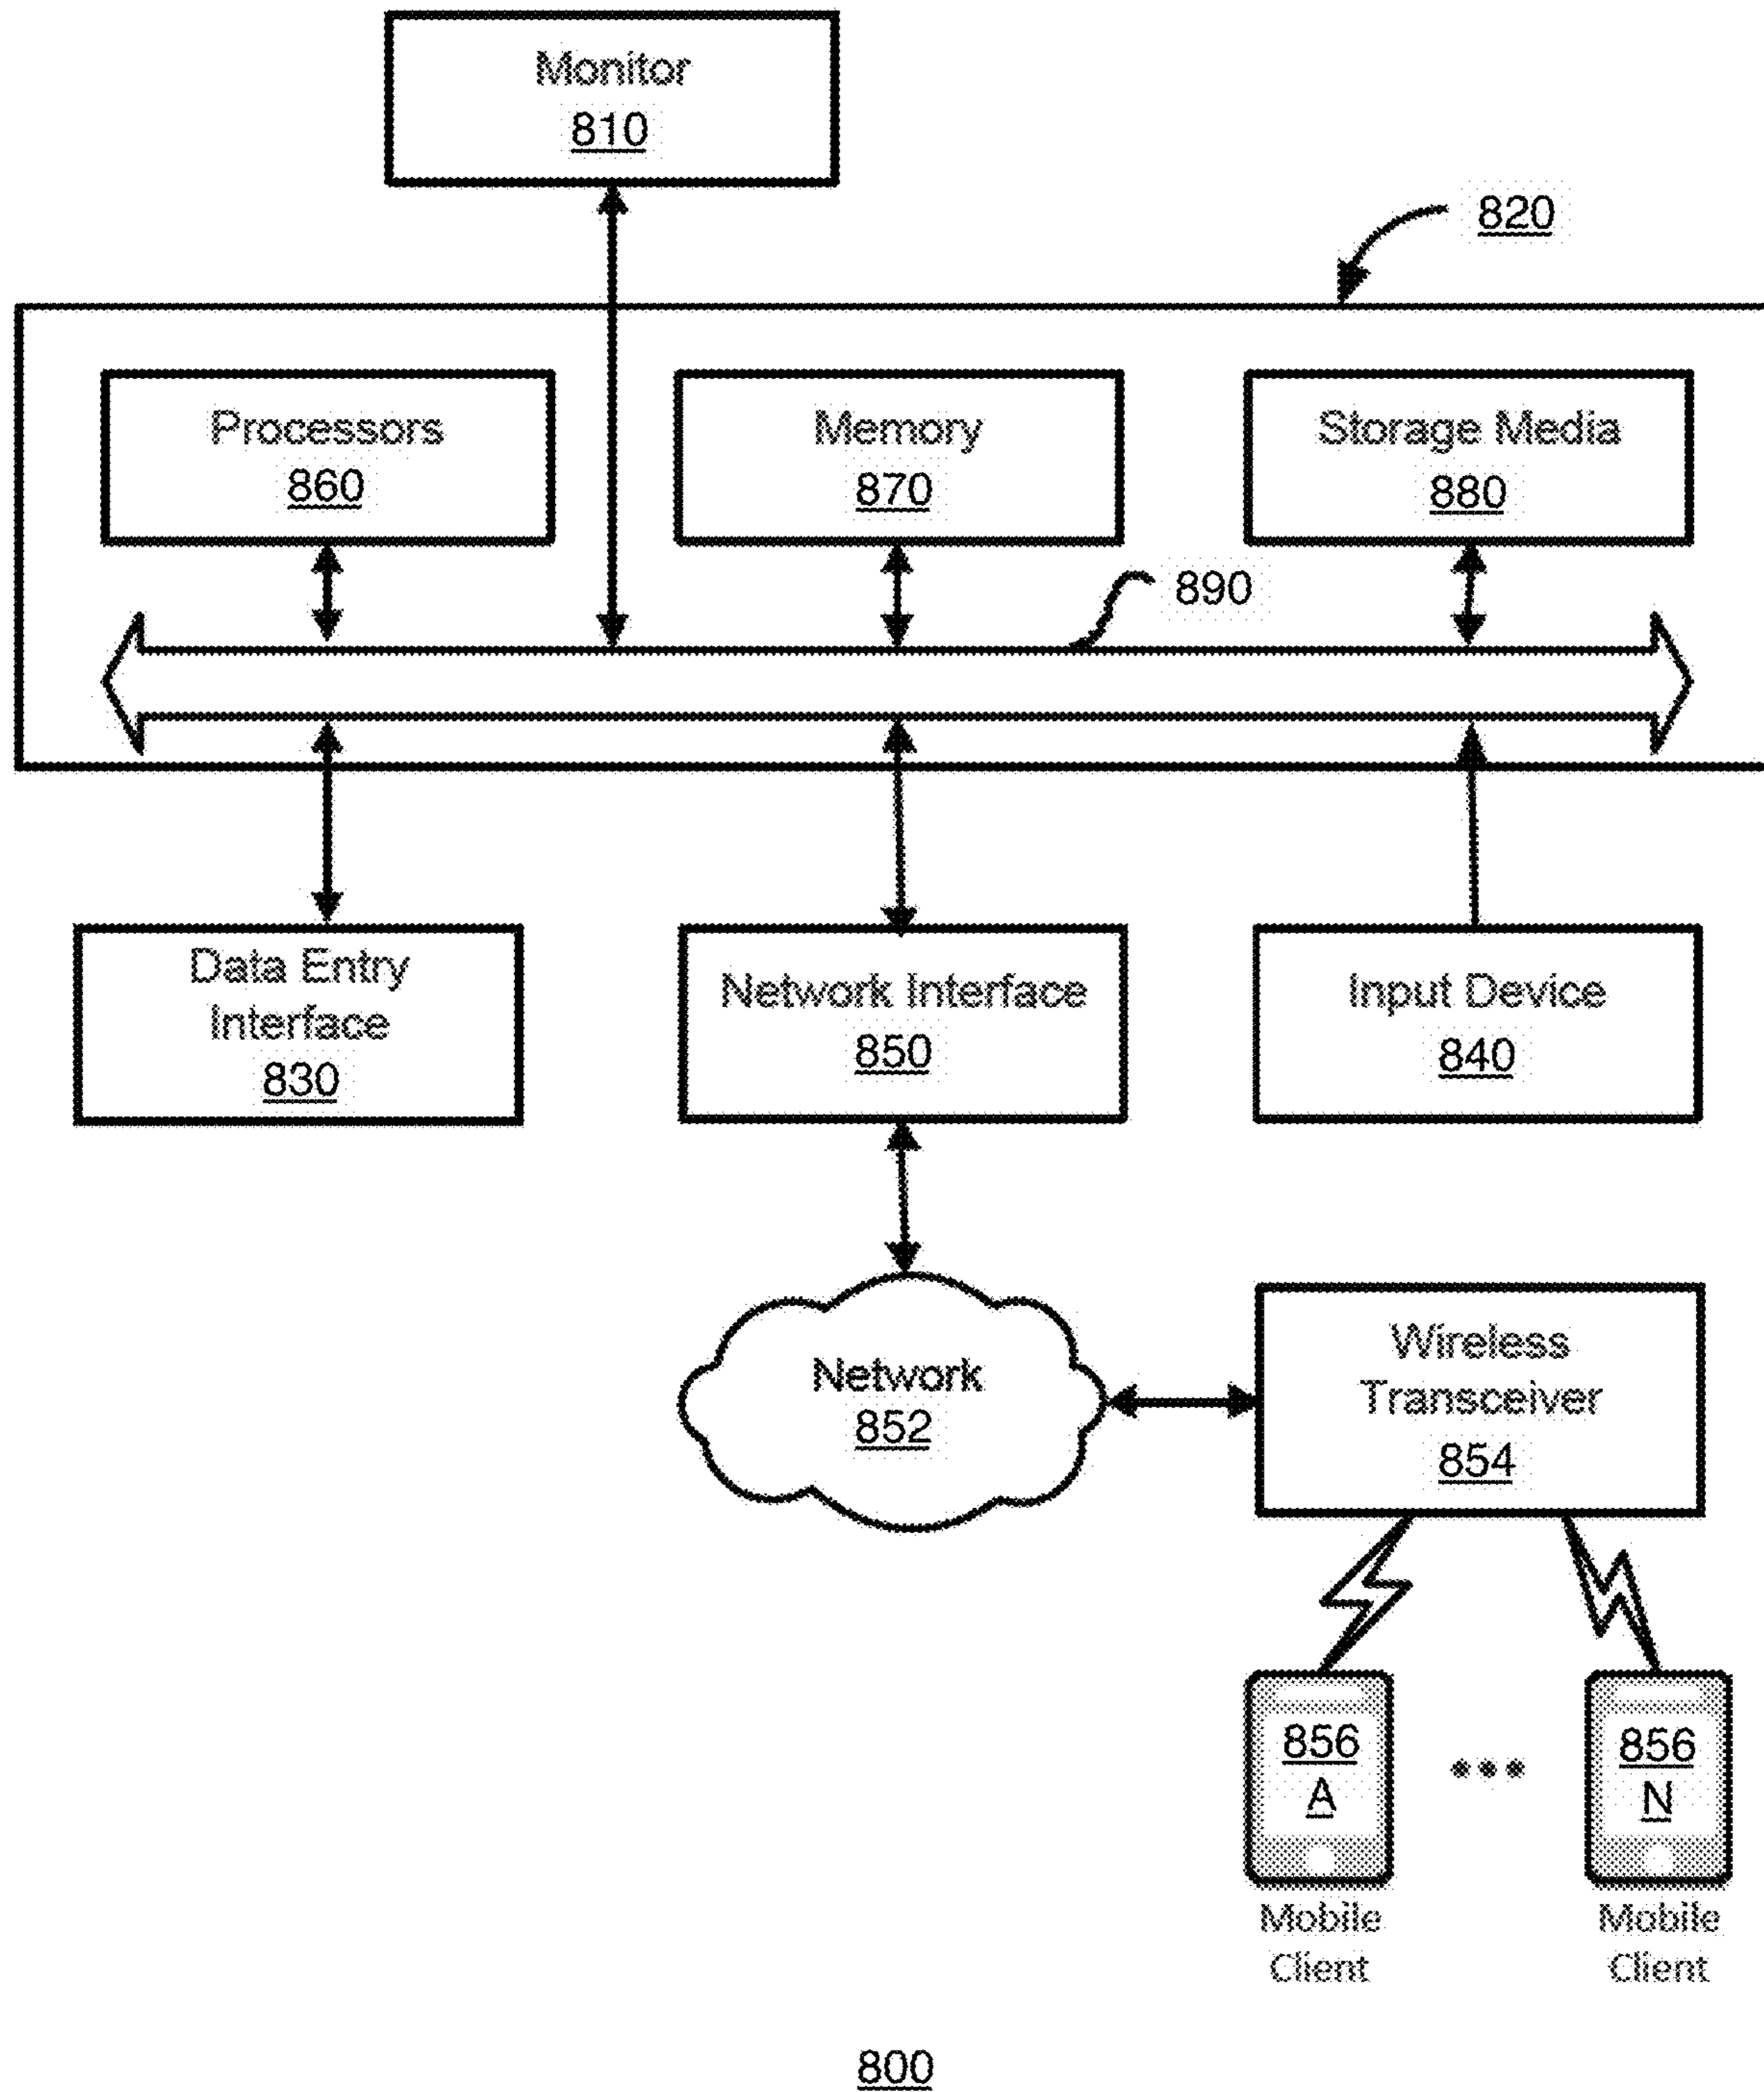


FIG. 8

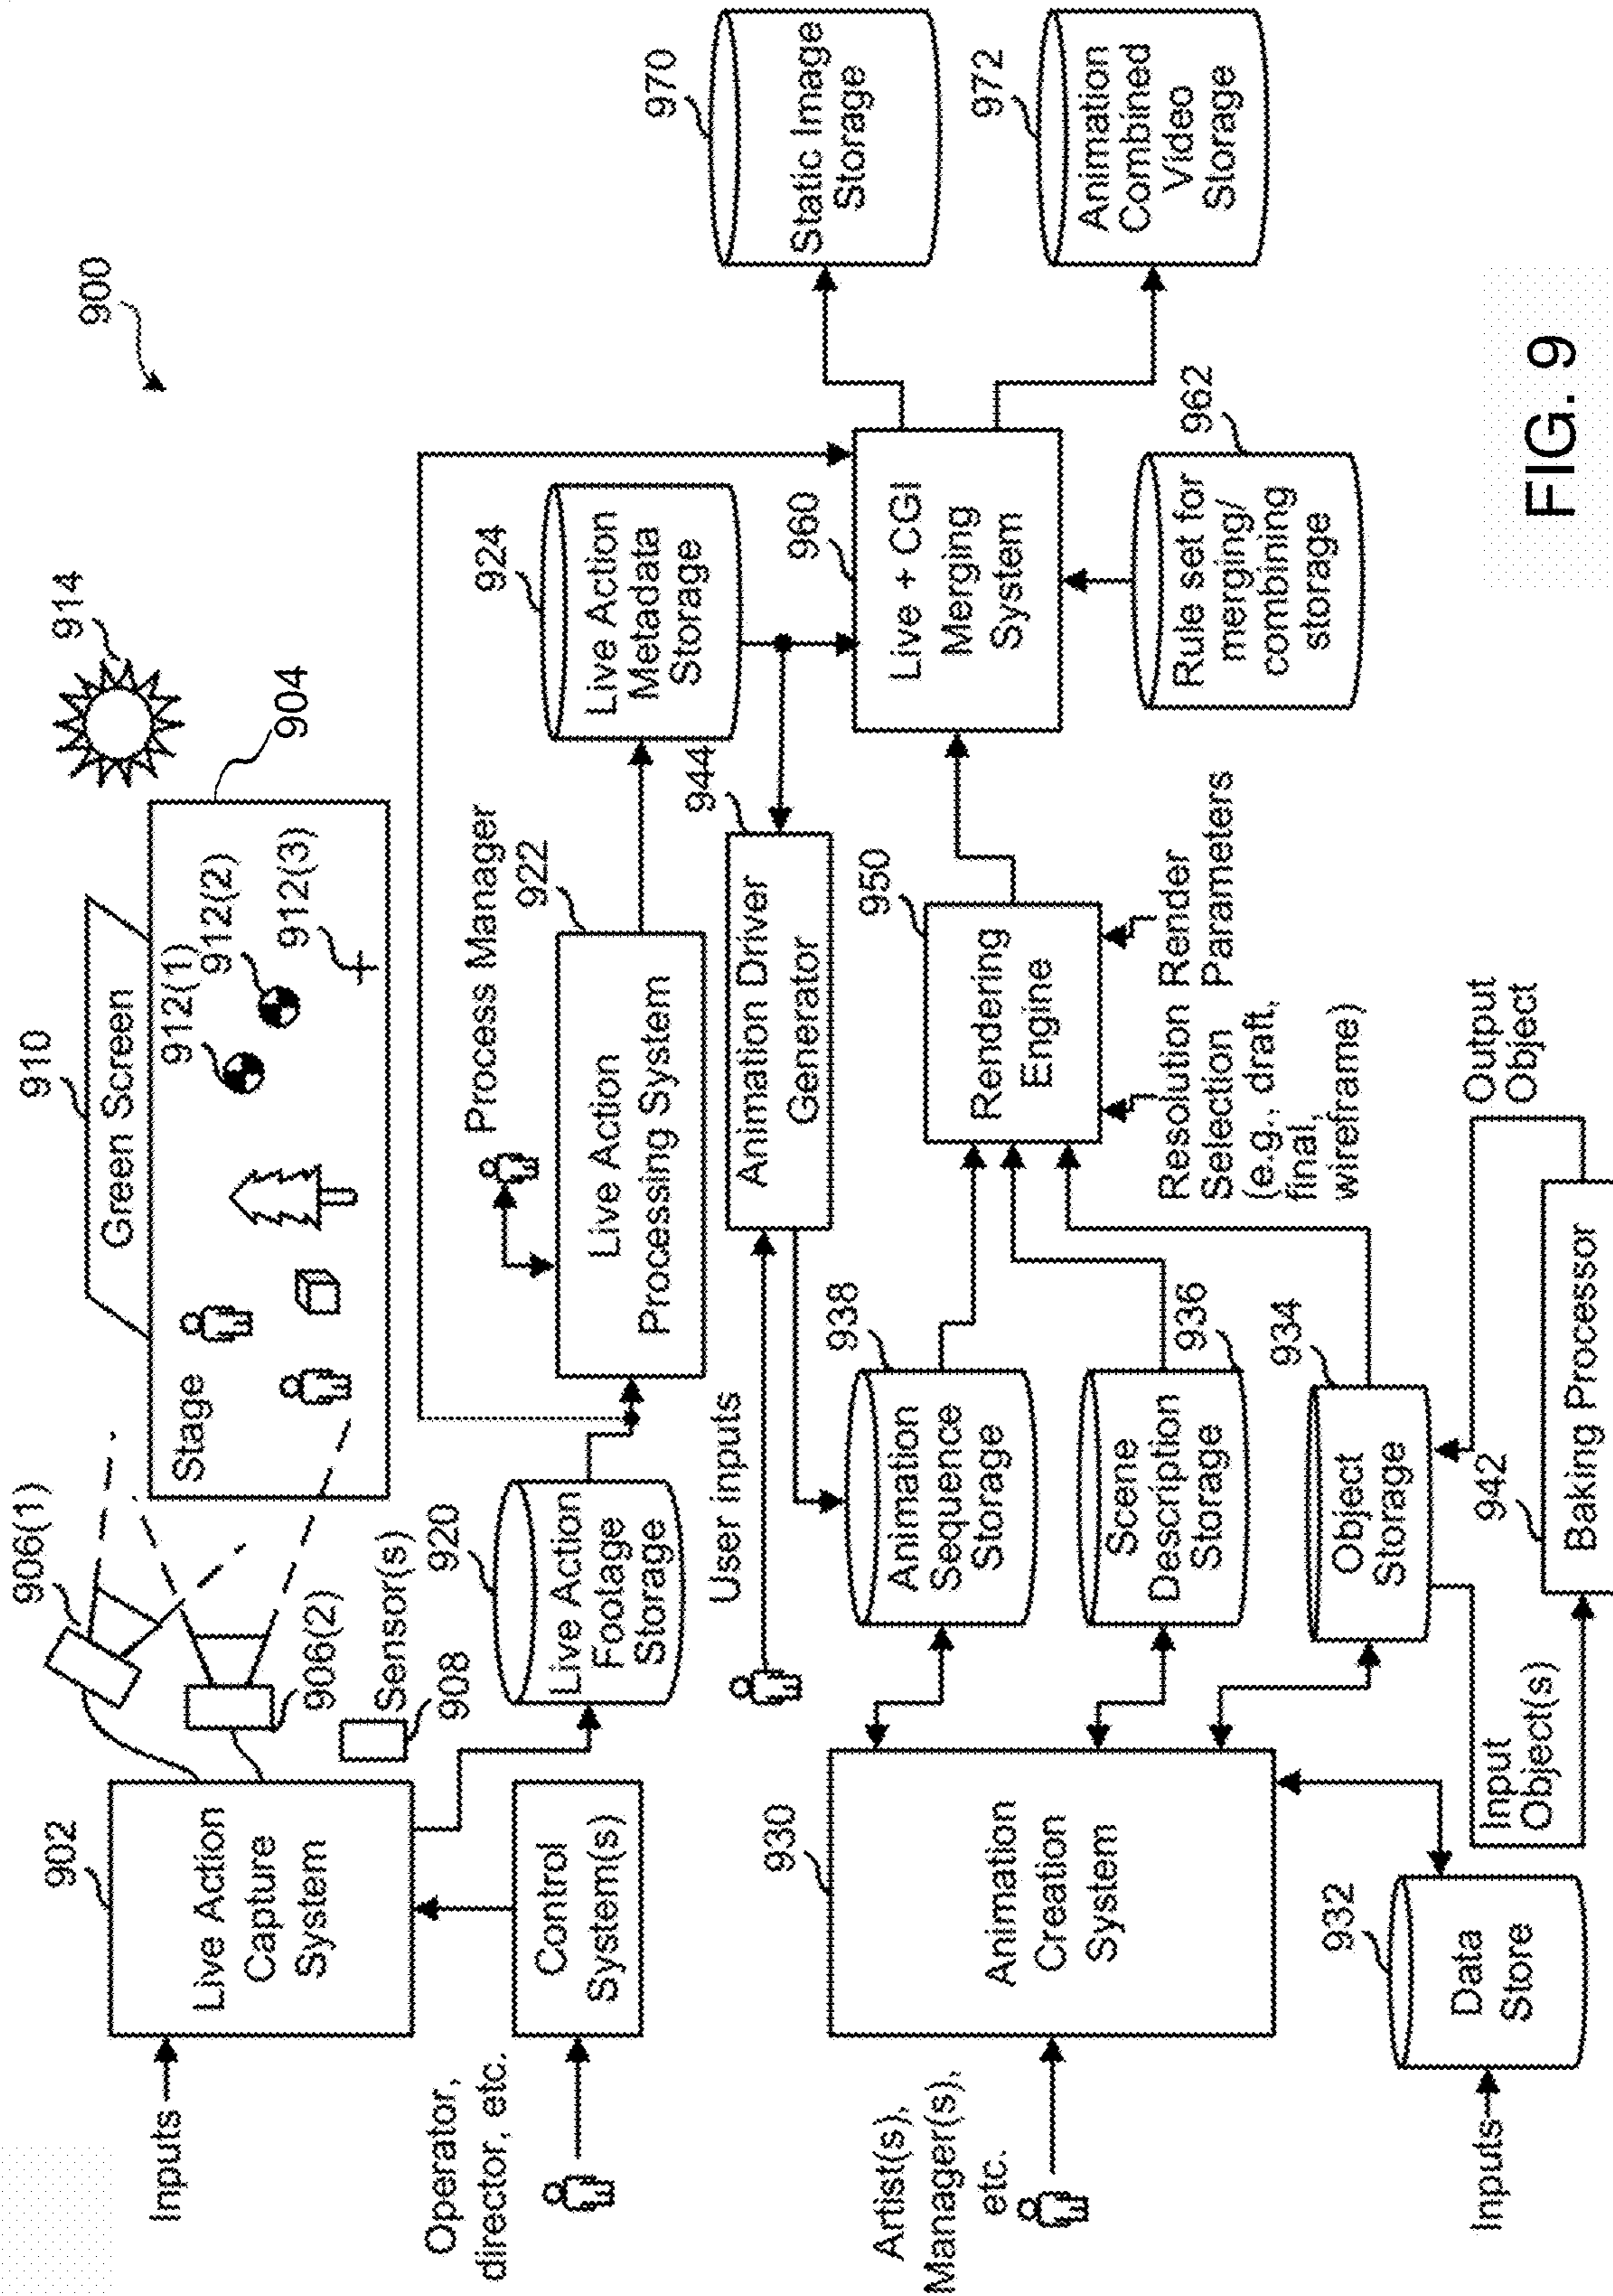



FIG. 9



1000  
FIG. 10

## DETECTION OF A CALIBRATION OBJECT FOR MODIFYING IMAGE PARAMETERS

### CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. patent application Ser. No. 17/065,439, entitled AUTOMATIC DETECTION OF A CALIBRATION OBJECT FOR MODIFYING IMAGE PARAMETERS, filed Oct. 7, 2020, which claims priority to U.S. Provisional Patent Application No. 63/047,868, entitled AUTOMATIC DETECTION OF A CALIBRATION OBJECT FOR MODIFYING IMAGE PARAMETERS, filed Jul. 2, 2020, which are hereby incorporated by reference as if set forth in full in this application for all purposes.

### BACKGROUND

Visual productions such as movies, videos, etc., require proper lighting for cameras to adequately capture people and objects in scenes. Even with proper lighting, the same object or objects captured in images may look different in different images depending on the cameras that captured the images. This may be due to image parameter values or operations varying from camera to camera. Variances in image parameters are typically corrected by manual adjustment of cameras or other equipment by artists or managers. Such adjustments may be imprecise, and cause delays in a visual production.

### SUMMARY

Implementations provide for automated detection of a calibration object within a recorded image, which may be used for various purposes such as modifying image parameters to make image capture more accurate and/or uniform. In some implementations, a system receives an original image from a camera, wherein the original image includes at least a portion of a calibration chart. The system further derives a working image from the original image. The system further determines regions in the working image, wherein the regions include groups of pixels having values within a predetermined criterion. The system further analyzes two or more of the regions to identify a candidate calibration chart in the working image. The system further identifies at least one region within the candidate calibration chart as a patch, wherein the identifying of the at least one region is based on a color of the patch. The system further predicts a location of one or more additional patches based on at least the identified patch.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of an example environment for detection of a calibration object in an image, which may be used for implementations described herein.

FIG. 2 is an example calibration chart, which may be used for implementations described herein.

FIG. 3 is an example flow diagram for detection of a calibration object in an image, according to some implementations.

FIG. 4 is an example calibration object, according to some implementations.

FIG. 5 is an example calibration object, according to some implementations.

FIG. 6 is an example calibration object, according to some implementations.

FIG. 7 is an example flow diagram for detection of a calibration object in an image, according to some implementations.

FIG. 8 shows basic components of an example computer system suitable for use with implementations described herein.

FIG. 9 is a block diagram of an example visual content generation system, which may be used to generate imagery in the form of still images and/or video sequences of images, according to some implementations.

FIG. 10 is a block diagram of an example computer system, which may be used for implementations described herein.

### DETAILED DESCRIPTION

Various implementations can replace or assist manual exposure and/or color correction processes. A calibration object used in a scene can be a calibration chart in the style of a so-called “Macbeth” chart, which may be a 4 by 6 grid of square patches. The particular dimensions of the grid may vary. In some implementations, six of the patches form a uniform gray lightness scale, and another six are primary colors typical of chemical photographic processes—red, green, blue, cyan, magenta, and yellow. The remaining colors may include approximations of medium light and medium dark human skin, blue sky, the front of a typical leaf, and a blue chicory flower. In some implementations, other patches may be chosen arbitrarily to represent a gamut of general interest and utility for test purposes. Orange and yellow patches may be similarly colored to typical oranges and lemons. As described in more detail herein, the system is designed to locate a calibration chart in an image. Although particular types of charts such as a Macbeth chart are described as examples, it should be apparent that other types of charts or objects may be used.

FIG. 1 is a block diagram of an example environment 100 for detection of a calibration object in an image, which may be used for implementations described herein. Shown in FIG. 1 is system 102, which receives images from camera 104. Also shown is a scene with objects that camera 104 captures in an image 106.

In various implementations, environment 100 may include any type and number of subjects, objects, lighting effects, etc. In an implementation, the image is a still image of a movie set to be used in a subsequent movie or video recording. In other implementations, features described herein may operate on a sequence of images such as a video. In other implementations, features may be performed in real time, or near real time, rather than substantially after a still image is captured and before live action shooting begins.

For ease of illustration, one camera is shown. In other implementations, camera 104 may represent any number of cameras, which may be of the same type of camera or may be of multiple types of cameras. Such cameras may include stand-alone cameras that can capture photos or video, dedicated video cameras, motion picture cameras, and/or other types of cameras. In other implementations, some or all processing steps performed by system 102 may be performed by one or more digital processors in, at, or proximate to the camera.

Image 106 is processed to identify various regions such as regions 108, 112, 114, 116, 118 and 122, in a scene or shot captured in environment 100. These regions are groupings of pixels of similar color and/or gray scale. As described in

more detail herein, system 102 may create the regions by grouping sets of adjacent pixels whose values fall within and/or meet predetermined criterion, or above or below threshold values. For example, region 108 on the teapot is an area of white represented by pixel values that vary because of the curvature of the teapot and the position of the light sources. A single region 108 is defined by using predetermined pixel value criterion and/or threshold settings. Similarly region 122 on the sugar bowl is identified based on other, or the same, criterion or threshold pixel values or other mathematical or logical criteria whether presently known or future-developed.

The region breakdown proceeds to identify many regions in the scene. There may be thousands or tens of thousands or more regions identified. The purpose of the region breakdown and some of the subsequent steps described below is to correctly identify the calibration chart 110 in the image by identifying color/gray scale patches such as patches 120 (e.g., color/gray scale patches in a Macbeth chart and the like). Ideally, enough regions of calibration chart 110 will be identified (and other regions rejected) to determine the correct position and orientation of chart 110 in image 106. Such identification and rejection can be difficult where patterns in image 106 are similar to calibration chart 110. For example, a plaid fabric pattern on the chairs can result in multiple regions 112 being identified in a regular array of different colors. Another example is regions 114, 116 and 118 on the table leg having different colors which, when identified as regions, may be later mistaken as color patches in a calibration chart. Many such visual similarities present obstacles to automated image detection including lighting effects (shadows, bright spots), noise, material reflective properties, etc.

As indicated above, at least one object captured in image 106 by camera 104 is a calibration object such as calibration chart 110. In some implementations, calibration chart 110 is a predetermined chart including known color patches 120. The color patches may be identified by their position in the chart regardless of environment (e.g., lighting) or camera settings. For example, it would be known that the patch in the lower-left corner of the chart is a light gray. In general, any organization of colors in known positions or relationships may be used. Alternatively, in some implementations, objects other than a “chart” may be used. For example, in some implementations, three-dimensional objects such as spheres, cubes, etc., may be used for calibrating scenes for three-dimensional capture such as in 3D or virtual reality recording. In various implementations, calibration chart 110 is placed in the scene for at least one captured image 106 of a series of images. The captured image that includes calibration chart 110 may be referred to as a source image in that the image provides a source of calibration information for corrections of one or more image parameters. The system may initially apply such corrections to the source image and ultimately to an entire series of related images that do not include a calibration. Because the related images are captured by the same camera, the correction being applied to all images appropriately corrects the image parameters in the same manner. Such image parameters may include brightness, white balancing, exposure, etc.

In various implementations, system 102 obtains image 106, which shows various objects, and identifies calibration chart 110 in the image. System 102 identifies patches 120 (enclosed by dotted lines) in calibration chart 110. As described in more detail herein, system 102 determines one or more colors from the patches. System 102 then determines at least one correction to at least one image parameter

associated with image 106 based on patches 120 in calibration chart 110. Further implementations directed to correcting image parameters of an image are described in more detail herein.

Also shown in FIG. 1 are regions 122, 124, 126, and 128 (indicated by or delineated by dashed lines) of image 106. In various implementations, system 102 generates the regions (e.g., regions 108, 112, 114, 116, 118, etc.) to analyze image 106, identifies calibration chart 110 and patches 120, and determines various position information associated with calibration chart 110 and patches 120.

For ease of illustration, 4 regions are shown. The number of regions that the system delineates in image 106 may vary, and the particular number of regions will depend on the particular implementation. Various implementations directed to regions 112, 114, 116, and 118 are described in more detail herein.

For ease of illustration, FIG. 1 shows one block for each of system 102, camera 104, and calibration chart 110. Blocks 102, 104, and 110 may represent multiple systems, cameras, and calibration charts. Also, there may be any number of objects captured in image 106. In other implementations, environment 100 may not have all of the components shown and/or may have other elements including other types of elements instead of, or in addition to, those shown herein.

FIG. 2 is an example calibration chart 200, which may be used for implementations described herein. In various implementations, calibration chart 200 may be used to implement calibration chart 110 of FIG. 1. In various implementations, calibration chart 200 may be a Macbeth color checker chart, or other suitable chart used for calibrating image parameters in one or more images. Further implementations involving calibration chart 200 and/or calibration chart 110 of FIG. 1 are described in more detail herein.

As shown, calibration chart 200 includes two sets of color patches, which include patches 202 and patches 204. Each set of color patches is indicated by dotted lines. In various implementations, patches 202 are red-blue-green (RGB) patches, and patches 204 are gray scale patches. The terms “color patches” and “patches” may be used interchangeably.

As shown, each patch in the sets of patches is unique in color and indexed with unique numbers. In various implementations, different patches of the calibration chart may have different colors. In various implementations, different patches may have the same color yet have varying shades of the same color. The various shades of a color may vary by different degrees of shading. The colors and index number may vary, and will depend on the particular implementation. In this example implementation, the upper left patch is a chromatic color of dark skin (labeled “Dark Skin”) and has an index number of “01.” The upper right patch is a chromatic color of bluish green (labeled “Bluish Green”) and has an index number of “06.” Other example chromatic colored patches are shown in the set of patches 202.

In another example, the lower left patch is a gray scale color white (labeled “White”) and has an index number of “19.” The lower right patch is a gray scale color black (labeled “Black”) and has an index number of “24.” Other example gray scale color patches are shown in the set of patches 204. While a total of 24 patches are shown, the particular number of patches may vary, and will depend on the particular implementation.

In various implementations, patches have a predetermined shape and a predetermined size. In various implementations, the predetermined shape of each of the patches is square.

The predetermined shape of the patches may vary, and will depend on the particular implementation.

In various implementations, the predetermined size of each of the patches may vary, and will depend on the particular implementation. In various implementations, the predetermined size is a minimum predetermined pixel number threshold. For example, in various implementations, the predetermined size may be a predetermined number of pixels in volume (e.g., 4 pixels per patch, 9 pixels per patch, etc.). This predetermined size or number of pixels may be referred to as a predetermined pixel number threshold.

As described in more detail below, in various implementations, the system analyzes the pixels of each patch to determine a particular color value (e.g., dark skin, light skin, blue sky, etc.) of each patch. In various implementations, the system requires that the number of pixels in a given patch meets at least the predetermined pixel number threshold in order for the system to adequately read the pixel values and determine the color of the pixels in the given patch.

FIG. 3 is an example flow diagram for detection of a calibration object in an image, according to some implementations. As described in more detail below, implementations identify image calibration data including information in the calibration object, which may be used to correct image parameters of an image. Referring to both FIGS. 1 and 3, a method is initiated at block 302, where a system such as system 102 receives an original image 106 from a camera, where original image 106 includes at least a portion of a calibration chart. In various implementations, image 106 includes objects, where at least one of the objects is a calibration chart 120, which may be captured completely or partially in image 106. Such images may be taken on location in the film setting and then may later be processed by system 102 or alternatively processed in real time on site. This may be a scenario where system 102 is on site and is communicating with camera 104 in real-time. In various implementations, image 106 may be referred to as a source image, where the source image is one image (e.g., the first image) of a series of images captured by the same camera.

At block 304, system 102 derives a working image from the original image. For ease of illustration, reference to image 106 of FIG. 1 may refer to the original image and/or to the working image unless otherwise specified. Because different cameras may each have a different color science with inherent variances of image parameters, different cameras may produce images with slightly different color characteristics. For example, a given camera may capture the same objects in a scene as another camera, yet both cameras have different levels of brightness. Such differences may be detectable by the human eye. As such, the brightness or other image parameters need to be calibrated or corrected, which is achieved with the various implementations described herein.

At block 306, system 102 determines regions in the working image, where each region comprises a group of pixels having values within a predetermined criterion. In various implementations, a given group of pixels may include one or more sets of pixels, where each set of pixels may contain connected, contiguous, and/or adjacent pixels. In various implementations, a given pixel of a set of connected pixels may be connected to 4 neighboring pixels, 8 neighboring pixels, etc., depending on the particular implementation. In various implementations, the predetermined criterion may include color and/or a color range. In some implementations, system 102 may search for a group of pixels that are within a predetermined distance from each other (e.g., within 1 pixel from each other, within 2 pixels

from each other, etc.). The predetermined distance may vary, depending on the particular implementation. For example, in various implementations, the predetermined distance may be that the pixels are adjacent to or neighboring each other.

As described in more detail herein, system 102 determines the regions in the image in order to identify calibration chart 110 in image 106. In some implementations, system 102 divides image 106 into regions 122, 124, 126, and 128, and searches regions 112, 114, 116, and 118 for patches 120 in calibration chart 110. In various scenarios, calibration chart 110 may be positioned within a single region. In other scenarios, calibration chart 110 may be span across multiple regions, where regions and/or portions of regions may be within calibration chart 110 in image 106. In various implementations, each region is a group of pixels with a particular color, and where the pixels are connected, contiguous, and/or adjacent pixels. As described in more detail herein, system 102 searches for and identifies calibration chart 110 in image 106 based on the identified regions.

At block 308, system 102 analyzes two or more of the regions to identify a candidate calibration chart in working image 106. As described in more detail herein, the patches have a predetermined shape and a predetermined size. As described in more detail herein, in some implementations, system 102 checks the shape of each region. In various implementations, system 102 may filter one or more portions of one or more regions based on shape. System 102 may filter or eliminate one or more portions of one or more regions and/or objects as potential or candidate patches 120 based on a predetermined shape. For example, of a portion of a region or an object is expected to be square or rectangular but system 102 determines it to be circular, system 102 may eliminate the region portion or object as a potential patch. As described in more detail herein, in some implementations, system 102 checks the size of each region. Also, in some implementations, the system may also eliminate one or portions of regions and/or more objects as potential or candidate patches 120 based on a predetermined size. Various example implementations directed to identifying calibration chart 110 and involving shape and size rejection are described in more detail below.

At block 310, system 102 identifies at least one region within the candidate calibration chart as a patch. In some implementations, the system determines a position of calibration chart 110 in image 106 based on the regions and based on the identified patches. As described in more detail herein, the system may identify one or more regions within the candidate calibration chart based on a color of the patch, including gray scale colors, as well as based on other factors (e.g., shape of patches in candidate calibration chart, patterns in the candidate calibration chart, etc.). As described in more detail herein, system 102 determines that various position information associated with calibration chart 110 may include fit and orientation of calibration chart 110 in image 106. Various example implementations directed to identifying patches 120 in calibration chart 110 are described in more detail below.

In various implementations, system 122 may determine one or more colors from patches 120. In various implementations, such colors may include chromatic colors and gray scale colors. Various example implementations directed to determining colors of patches 120 in calibration chart 110 are described in more detail below.

In various implementations, the system identifies color/gray scale patches such as patches 120 (e.g., color/grey scale patches in a Macbeth chart and the like) shown on calibra-

tion chart **110** in FIG. **1**, patches **202** and **204** shown on calibration chart **200** in FIG. **2**, etc.

The system detects different patches having variations of shade of colors, including gray scale colors. For example, the system may detect not only black and white shapes but also shapes having different degrees of shading, such as the patches shown in FIG. **2**.

In various implementations, the system also detects different patches at different orientations of the calibration chart relative to the cameras in the scene. The system detects various types of real objects in a photograph even when such objects are not square to the camera or in sharp focus. For example, the system may detect patches at glancing angles and/or obscured by reflections.

Implementations described herein are advantageous over some objects such as matrix or two-dimensional barcodes such as quick response (QR) codes, etc. This is because barcodes such as QR codes have many squares to decode and the squares are all black and white. In some scenarios, a barcode might not be facing a given camera squarely or completely. Less-than-robust techniques would be able to detect black and white patterns of a QR chart if adequately facing a camera squarely. However, less-than-robust techniques would not be able to detect black and white patterns of a QR chart if not positioned sufficiently squarely to a camera. Also, such techniques would not be able to identify similar shades of color on the calibration chart.

In some implementations, the system may detect a chart that is projected onto a surface. For example, a transparency slide in a projector may contain patches that reflect different wavelengths of light. If photographed, the system may analyze the projected calibration chart to decode the spectral response of at different patches detected on the surface.

In addition to detecting calibration charts having various numbers, shapes, and colors and shading of patches, the system may also detect patches in order to profile sensor noise, camera focus, and lens sharpness of cameras.

At block **312**, system **102** predicts a location of one or more additional patches based on at least the identified patches. For example, after identifying at least one patch, system **102** may then analyze other regions in image **106** to identify other patches in a similar manner. In various implementations, the contents of calibration chart **110** is known or predetermined. As such, system **102** computes or ascertains the location of identified patches captured in image **106**.

In various implementations, in response to the one or more colors of the patches **102**, system **102** determines at least one correction to at least one image parameter associated with image **106** based on patches **120** in calibration chart **110** as further described below with reference to FIG. **4**. For example, referring to FIGS. **1** and **2**, system **102** may identify a patch of chart **120** as “21 Neutral 6.5” and utilize a reference RGB value of patch “21 Neutral 6.5” to determine the correction value to apply to RGB used to generate image **110** in order to change the determined value of “21 Neutral 6.5” in image **110** to within a tolerance of the reference RGB value. In various implementations, system **102** may determine a multiplier to calibrate or correct the image parameter of image **106**. System **122** may then apply the multiplier to all images of a series of images of a particular video footage.

In various implementations, system **102** generates one or more image parameters, where one of the image parameters may include a brightness parameter. In various implementations, system **102** generates one or more image parameters, where one of the image parameters may include a exposure

parameter. The particular image parameter that system **122** adjusts and the number of image parameters that system **122** adjusts may vary and will depend on the particular implementation. For example, in various implementations, system **102** generates one or more image parameters, where one of the image parameters may include a white balance parameter. In other implementations, the image parameter may be hue. Various example implementations directed to determining one or more corrections to one or more image parameters are described in more detail below.

Although the steps, operations, or computations may be presented in a specific order, the order may be changed in particular implementations. Other orderings of the steps are possible, depending on the particular implementation. In some particular implementations, multiple steps shown as sequential in this specification may be performed at the same time. Also, some implementations may not have all of the steps shown and/or may have other steps instead of, or in addition to, those shown herein.

FIG. **4** is an example calibration object **400**, according to some implementations. As shown, calibration object **400** is a cube. In various implementations, one or more color charts and/or color patches may be positioned upon calibration object **400**. For example, calibration object **400** includes calibration charts on different faces such as calibration charts **402**, **404** and **406**. Some of the other three faces (not shown) may also have calibration charts.

In various implementations, the steps and analysis can proceed similarly to those discussed herein for a two-dimensional calibration chart. An advantage of calibration object **400** is that it provides multiple samples of the same color patch (such as color patches **412**, **414** and **416** which are identically colored) exposed to different lighting planes in the scene. This provides additional pixel values and modes for analysis to obtain parameter information for camera calibrations, digital rendering, modeling, etc.

While calibration object **400** is shown in this example as a cube, in other implementations, a calibration object may be other types of solid objects of different shapes and sizes (e.g., a sphere or other curved surface). In some implementations, the object may simply be uniformly colored with a single solid color. Other coloring approaches such as stripes that cut across planes of the object, etc., may be used. Arrays of multiple objects can be provided.

Because multiple faces are now presented to a given camera, the system may detect the effects of multiple environment light sources, reflections, filters, and/or other effects by the way in which the colors on the object are captured by the camera and/or other cameras. For example, if there is an incandescent light source behind a given camera, and there is also an overhead fluorescent light source, the recorded color of patches **410**, **412**, and **414** will vary. This presumes the color charts on each face/side of the cube are the same. In other words, while patches **410**, **412**, and **414** are the same color, the system will record the color of these patches differently, depending on the lighting conditions.

In various implementations, each of the color pixel values correspond to the same color or shade of gray in the scene and the color can thus be matched even under the different lighting conditions that exist within the scene. The color object with multiple color patches and/or charts can convey differences in calibration dependent upon camera and subject position and orientation.

To further illustrate, in a similar scenario to that described above, presume that there is an incandescent source behind a given camera and an overhead fluorescent light source.



Then, if the camera changes position (e.g., raises or flies above the calibration object), the calibration for the colors would be changing in the different color temperatures of the different lights. The use of data from the calibration object can provide dynamic calibrating or correction during scenes where the camera, subjects, and lights are moving in the scene.

Still referring to FIG. 4, the normal of each face of the cube are shown as arrows pointing out from each of the three visible faces. Because these are normal to each other, the system may interpolate the calibration information to derive the calibration effects for any orientation in the visible hemisphere (e.g., front-facing, etc.) portion of an object in the scene.

The information from calibration object 400 can be useful when integrating or compositing digital effects within a filmed sequence. This is because any digital items can be made to reflect or visually react in the same manner as the physical items. Calibration objects may also be useful when recording 3D sequences, and/or in virtual reality, and/or augmented reality, and/or in other three-dimensional presentations.

FIG. 5 is an example calibration object 500, according to some implementations. As shown, calibration object 500 is a quad-card type object where four calibration charts are designed onto a four-sectioned placard. In some implementations, markers such as 502 may be used to enable a camera or other equipment to determine the orientation and position of the quad-card in space in the scene. To determine angles of the quad-card's sections, the sections may be positioned by a human operator in order to suit the lighting and effects conditions of the scene. Any means for detecting the position and orientation of the sections may be used such as image processing, sonar, radar, lidar, reflective or active markers, etc.

FIG. 6 is an example calibration object 600, according to some implementations. As shown, calibration object 600 is a dodecahedron having 12 sides or faces (only 6 faces shown). Also, each face has a calibration chart positioned thereon. Such a calibration object is useful where it is desired to have at least one of the calibration charts be substantially facing the camera. This substantially replicates characteristics and benefits of a single calibration chart approach, as described in FIG. 2, for example. Although the faces of the dodecahedron are not normal with respect to each other, the system may still compute interpolations for calibration. In some implementations, not all of the faces (e.g., rear-faces, etc.) of such a calibration object need to be used. For example, the rear of calibration object 600 may be a flat plane or other shape in order to conserve material and space. Furthermore, in some implementations, this enables the calibration object to be positioned or mounted on different surfaces or by different equipment. For example, a boom attachment socket may be incorporated in the back of the object.

FIG. 7 is an example flow diagram for detection of a calibration object in an image, according to some implementations. Referring to both FIGS. 1 and 7, a method is initiated at block 702, where a system such as system 102 obtains image 106. In various implementations, image 106 includes objects, where at least one of the objects is calibration chart 110. As indicated above, in various implementations, image 106 may be referred to as a source image, where the source image is at least one image (e.g., the first image) of a series of images captured by same camera 104. As described in more detail below, the system may then make the correction to the image parameter source image

and then apply the same correction to the other images of the sequence. Applying the same adjustment for all images of the series will have correct results. Because the images of the series are captured by the same camera, any correction will fix any variance of an image parameter such as brightness caused by the camera.

At block 704, system 102 adjusts the resolution of image 106. In some implementations, the system may adjust (e.g., reduce) the resolution of image 106 downward in order to increase the processing speed and reduce needed computation resources.

At block 706, system 102 determines regions 112, 114, 116, and 118 in image 106. In some implementations, the system divides image 106 into regions 112, 114, 116, and 118. The number of regions that the system determines or delineates in image 106 may vary, and the particular number of regions will depend on the particular implementation. As described in more detail herein, in various implementations, the system searches on or more of regions 112, 114, 116, and 118 for patches 120 in calibration chart 110. As described in more detail herein, in various implementations, system 102 determines a position of calibration chart 110 in image 106 and its associated patches 120 based on the regions.

Note that the determining of regions 112, 114, 116, and 118 is described above in connection with FIG. 2 in the context of the system identifying calibration chart 110 in image 106. In various implementations, the step of determining regions may be performed before and/or as a part of the system identifying calibration chart 110. In some implementations, the step of block 406 may be optional.

In some implementations, system 102 may utilize a region-merging algorithm to group together pixels that have the same color. The system does not make assumptions about the color of the patches of calibration chart 110 in image 106, but rather reads the actual image parameter values from the pixels in image 106. As described in more detail herein, the system compares these actual, read image parameter values to predetermined correct or desired image parameter values in order to determine appropriate corrections. For example, the system may correct a patch that is gray but that is not sufficiently bright by applying a multiplier to increase the brightness to the desired brightness value. The correction will correct all other pixels of the same gray color in the region in the same corrective manner.

The step of block 706 is beneficial in that it enables system 122 to more accurately and robustly identify calibration chart 110. For example, by dividing or delineating image 106 into regions, system 122 systematically searches for and identifies calibration chart 110, and its associated patches 120, region by region. In various implementations, once calibration chart 110 and its associated patches 120 are found and identified, system 122 may stop analyzing image 106, unless system 122 needs to search for and find a second calibration chart and associated patches. By being able to analyze portions of image 106 instead of potentially having to analyze entire image 106, system 102 consumes fewer computer resources. This also saves computation time by not needing to take extra time to further process regions that might not contain a calibration chart 110. Furthermore, analyzing image 106 region-by-region enables system 122 to search and identify calibration chart 110 and its associated patches 120 faster by potentially eliminating unnecessary processing of objects on at least some regions.

At block 708, system 102 performs shape rejection. As indicated above, the patches of calibration chart 110 have a predetermined shape. In various implementations, system 122 may eliminate one or more objects as patches 120 based

## 11

on the predetermined shape. In other words, the different objects captured in image 106 are candidate objects, where system 122 determines which objects of the candidate objects are patches of calibration chart 110. As system 122 passes through regions of image 106, system 122 compares each object to the predetermined shape and rejects or eliminates candidate objects that do not match the predetermined shape. For example, in various implementations, the predetermined shape is a square. As such, system 122 will reject a rounded, oblong object as a patch. In some implementations, the system 122 may compensate for angle and perspective in searching for objects that may vary in shape relative to their angle to an image plane of camera 104, such as squares. In some implementations, system 102 may analyze the aspect ratio of each object. The system may reject or eliminate candidate objects where the height and width are not equal. In this example, non-square objects (e.g., a person, cylindrical barrel, etc.) would not be patches. Conversely, the system may identify candidate objects as potential patches if such candidate objects have equal height and width.

In some scenarios, a particular patch may be split into two portions due to shadows, noise, dirt, etc. In some implementations, the system may determine that two portions are a part of the same patch based on the color read from each of the patch portions.

At block 710, system 102 performs size rejection. As indicated above, the patches of calibration chart 110 have a predetermined size. In various implementations, the system may eliminate one or more candidate objects as patches based on the predetermined size. As the system passes through regions of image 106, the system measures each object and rejects or eliminates candidate objects that are smaller than the predetermined size. As indicated above, for example, in various implementations, the predetermined size may be a predetermined number of pixels in volume (e.g., 4 pixels, 9 pixels, etc.). This predetermined number of pixels may also be referred to as a predetermined pixel number threshold. In some implementations, system 102 may analyze the number of pixels of each object. System 102 may reject or eliminate candidate objects where the number of pixels falls below the predetermined pixel number threshold. In this example, objects that are too small (e.g., under 4 pixels) would be rejected as patches. Even if such objects were indeed patches, the system may determine that the patches may be too small for analysis. For example, a given calibration chart 110 may be located too far back in the background of the image, and would be too small. If a patch is too small, system 102 might not appropriately or accurately read and determine color information from the patch. As such, system 102 might not recognize the object as a patch 120. Conversely, system 102 may identify candidate objects as potential patches 120 if such candidate objects meet or exceed the predetermined pixel number threshold.

In various implementations, system 102 may utilize suitable pattern recognition techniques to recognize and identify candidate objects as patches for such objects that meet both the predetermined shape and predetermined size requirements. In some implementations, system 122 may recognize and identify candidate objects as patches for such objects that meet at least one of the predetermined shape and predetermined size requirements.

As indicated herein, system 102 searches for a complete set of patches (e.g., 18 chromatic patches, 6 gray scale, etc.) or the most complete set of patches. In various implementations, system 122 need only find a subset of the total number of patches order to identify calibration chart 110.

## 12

For example, if a patch is black and the background is dark or otherwise obscured, system 102 might not find that particular black patch. In some implementations, system 102 may search of a portion (e.g., 4, 5, or 6 patches) of a row. As long as one patch is identified, system 102 will ascertain the location of calibration chart 110.

In various implementations, system 102 computes the orientation of the candidate calibration chart. If more than one patch is identified, it increases the robustness of system 102 ascertaining the location of calibration chart 110, as well as its orientation. As such, the system 102 may identify only a portion of all patches 120 of given calibration chart 110. Implementations described herein may still operate successfully based on the patches that are identified, depending on the particular implementation. For example, one or more patches 120 may be obscured by another object, or visual “noise”, or a shadow, etc. As such, system 102 might not identify all patches 120 of a given calibration chart 110. In various implementations, as long as system 102 identifies certain patches 120, the system 102 may still read the colors of the identified patches 120 and still perform the various steps to correct one or more image parameters of image 106.

Note that while the performance of shape rejection and size rejection is described above in connection with FIG. 2 in the context of system 102 identifying calibration chart 110 in image 106, in various implementations, the steps of shape rejection and size rejection may be performed before and/or as a part of the system identifying calibration chart 110. In some implementations, the steps of blocks 708 and 710 are optional, where the system may utilize both, either one of, or none of the steps of blocks 708 and 710.

The steps of block 708 and/or block 710 are beneficial in that they enable system 102 to more accurately and robustly identify calibration chart 110 by eliminating false positives (e.g., eliminating a non-patch object from candidate patch objects. For example, by eliminating all or at least some objects in image 106 (e.g., objects 108, 112, 114, 116, etc.) as candidate objects to be determined as patches of calibration chart 110, system 102 consumes fewer computer resources by avoiding further analysis on such objects. This also saves computation time by not needing to take extra time to further process non-patch objects. Furthermore, the process of eliminating non-patch objects from candidate objects enables system 102 to search and identify actual patches of calibration chart 110 faster by eliminating unnecessary processing of the non-patch objects. Furthermore, the process of eliminating non-patch objects from candidate objects enables system 102 to more accurately identify patches of calibration chart 110.

At block 712, system 102 identifies calibration chart 110 in the image. In various implementations, system 102 identifies the calibration chart in the image based at least in part on the system’s performing of shape rejection and size rejection. System 102 finds and identifies patches 120 based on shape and size of patches 120, and system 102 eliminates non-patch objects based on shape and size of those objects. By identifying patches 120, system 102 identifies calibration chart 110, which contains patches 120.

At block 714, system 102 determines if calibration chart 110 is of a sufficient size. If yes, system 102 continues to block 718, described below. If no, system 102 continues to block 716, described below. The calibration chart might not be of a sufficient size in a scenario, for example, where calibration chart 110 is small in the image (e.g., far away).

At block 716, system 102 crops calibration chart 110 from the original image 106. After cropping calibration chart 110, system 102 returns to block 704, described above.

At block 718, system 102 determines the fit of calibration chart 110 in image 106. As indicated above, in various implementations, system 102 determines a position of calibration chart 110 in image 106 and its associated patches 120 based on the regions. One aspect of the position is the fit of calibration chart 110, and another aspect of the position is the orientation of calibration chart 110, which is described below in connection with block 716.

With regard to the fit of calibration chart 110, in various implementations, system 102 determines the coordinates of calibration chart 110. For example, system 102 may determine that the upper left corner of calibration chart 110 is at a coordinate at 0,0, the upper right is at 1,0, the lower left is at 0,1, and the lower right is at 1,1. System 102 may then perform a mathematical transform to convert the calibration chart space coordinates to pixel coordinates. As noted herein, patches 120 of calibration chart 110 are located in predetermined positions of calibration chart 110. As such, once system 102 determines the corners of chart 110, system 102 ascertains the locations of patches 120 in the image. Furthermore, because system 102 knows the geometry of calibration chart 110 and its patches 120, system 102 ascertains a location, such as the center, of each given patch 120 for reading image parameters such as brightness, etc.

At block 720, system 102 determines the orientation of calibration chart 110 in image 106. In various implementations, system 102 determines the orientation of calibration chart 110 based on determining the relative coordinates of at least 2 patches. Any 2 identified patches in image 106 may inform the orientation of calibration chart 110. For example, the relative positioning of the 24 patches of calibration chart 110 is known. As such, if the system determines that an upper left patch is "Dark Skin 01" and the lower right patch is "Black 24," system 102 may ascertain that calibration chart 110 is right side up. Conversely, if the system determines that an upper left patch is "Black 24" and the lower right patch is "Dark Skin 01," system 102 may ascertain that calibration chart 110 is upside down. Every additional patch 120 that is identified incrementally improves the accuracy and robustness of identifying the orientation of calibration chart 110.

Note that while the determining of the fit and orientation of calibration chart 110 is described above in connection with FIG. 2 in the context of the system identifying patches of calibration chart 110, in various implementations, the steps of determining the fit and the orientation may be performed before and/or as a part of system 102 identifying the patches of calibration chart 110. In some implementations, the steps of blocks 718 and 720 are optional, where system 102 may utilize both or either one of or none of the steps of blocks 718 and 720.

The steps of block 718 and/or block 720 are beneficial in that they enable system 102 to more accurately and robustly identify individual patches 120 of calibration chart 110. For example, by determining the fit and orientation of calibration chart 110, once system 102 determines one or more patches 120 and knows the color information from those patches, system 102 may quickly identify the rest of patches 120 in calibration chart 110, which may save computation resources and computation time.

At block 722, system 102 identifies patches 120 in the calibration chart 110. In various implementations, system 102 may utilize any suitable pattern recognition techniques to identify the patches based on shape and size. System 102 improves the robustness of identifying patches by eliminating the non-patch objects based on shape and size rejection techniques. In some implementations, when patch 120 is

identified, the system may enable the patch to indicate vote as to the configuration of chart 110, similar to a Hough transform.

At block 724, system 102 determines one or more colors from patches 120. In various implementations, system 102 computes or ascertains the orientation of the calibration chart based on the determined colors from patches 120. By knowing the calibration chart orientation and knowing the approximate centers or other location within patches 120, system 102 may then read the RGB pixel values in each patch of the image to determine the colors. System 102 may also read other image parameters such as brightness during the reading process. In various implementations, system 102 may increase the image resolution for color extraction. In various implementations, there may be various reasons to locate a chart and to extract recorded red-blue-green (RGB) values for each patch. For example, a reason may be to find adjustments to camera settings or post-processing to produce the required image.

In some implementations, in addition to color, shape, and other patch characteristics, the system may utilize patterns to enhance detection of a calibration chart, including subsequent camera calibration. For example, in some implementations, a given calibration chart may include a pattern of particular patches. These patches may be arranged in a particular color pattern (e.g., lightest to darkest, etc.). In some implementations, patches may be arranged in a particular shape pattern of patches having different shapes (e.g., polygons having fewer corners to polygons having more corners, etc.). In some implementations, patches may be arranged in a particular shape pattern of patches that have different sizes (e.g., smaller to larger, etc.).

In some implementations, in lieu of or in addition to patches of a calibration chart having one or more pattern aspects, a given calibration chart may include a predetermined pattern such as a large-code style. This predetermined pattern may be positioned at various locations on the calibration chart (e.g., center, top, bottom, side, etc.), and patches may be arranged around the predetermined pattern. For example, in some implementations, the patches may be arranged in an hexagonal arrangement surrounding the predetermined pattern. Such an arrangement enables the patches to be positioned closer to the center pattern, which would make locating them more reliable.

In some implementations, placing the predetermined pattern in the center of the calibration chart is beneficial in that objects such as user hands are less likely to obscure the predetermined image. For example, if a user is holding a calibration chart while the chart is being captured by a camera, the user is more likely to hold the chart at a side rather than at the middle of the chart. When the system knows the position and orientation of the predetermined pattern, the system may ascertain the locations and positions of the other patches in the image, regardless of the position of the calibration chart.

In some implementations, the system may also derive geometric information from a calibration chart. For example, the system may compute the relative orientation between the camera and the chart. Also, from knowledge of information about the lens, the system may ascertain the distance between the calibration chart and the camera. As such, as a user walks around the set with a calibration chart, the system may gather lighting information from the different positions and orientations.

At block 726, system 102 determines at least one correction to at least one image parameter associated with the image based on the patches in calibration chart 110. In

various implementations, system 102 extracts pixel colors from patches 120 of chart 110. System 102 then compares each read or actual pixel value to a target or reference pixel value to determine a difference. The reference pixel value may be derived from a reference calibration chart. Patches 120 should be identical, or within a tolerance, in color and brightness relative to a reference chart value. In various implementations, the difference may become an offset value or correction value. In various implementations, the output (e.g., correction value) may be stored in a spreadsheet, data structure, or other suitable file for further processing.

As indicated above, there may be various reasons to locate a chart and to extract recorded RGB values for each patch. In another example, a reason may be for system 102 to process the recorded RGB values for one or more cameras and/or any digital post-processing any system to perform. In this scenario, in some implementations, instead of performing the step of block 726, system 102 may compute and record average RGB color values and/or other statistics for all patches to be processed and/or analyzed by some other system.

As indicated above, in various implementations, system 102 may determine a multiplier to calibrate or correct the image parameter of image 106. More specifically, a calibration or correction to an image parameter means system 102 adjusts an actual image parameter value to a correct or desired image parameter value. For example, in some implementations, the system may apply the multiplier to the brightness value of each of the RGB channels. The system may then apply the multiplier to the brightness value of all images of a series of images of a particular video footage captured by same camera 104. With such adjustments, patches 120 on calibration chart 110 will have the correct or desired chromatic and/or gray scale brightness value and/or other correct image parameter value(s). As indicated above, system 102 may apply any correction to the source image to other images captured by the same camera to get similar or substantially the same results such as appropriate brightness or other image parameters (e.g., exposure, etc.). In various implementations, system 102 may use any corrections to discover information about a given camera, where such information may be unknown. In some implementations, system 102 may use such outputs or corrections for correcting camera settings.

In some implementations, the image parameter is brightness. In some implementations, the image parameter is exposure. The particular image parameter that system 102 adjusts and the number of image parameters that system 102 adjusts may vary and will depend on the particular implementation. Various example implementations directed to determining one or more corrections to one or more image parameters are described in more detail below.

In various implementations, the system measures a dynamic range of the camera based on one or more of a reflection and a illumination of the patch. In various implementations, a given calibration chart may include one or more patches having reflective and/or illuminated patches. For example, some patches may have different reflective properties, where the patches have a range of reflective values. In another example, some patches may have different illumination properties, where the patches have a range of illumination values.

The system may then compare reflective values and/or illumination values against corresponding reference values. Such reference values may include predetermined benchmark reflective and/or illumination values. Such reference

values may also include reflective and/or illumination values associated with other cameras.

In various implementations, the system detects multiple patches with different colors to facilitate in reverse engineering the dynamic range of a camera including its response to different colors. Adding some patches having different reflective and/or illumination properties facilitates the system in deriving more subtle information about the relationship between the spectral emission of each light and the spectral sensitivity of the camera. In some scenarios, if a given calibration chart has multiple grey patches, the patches would look the same under spectrally uniform light. However, under different types of light, the system detects different colors in patches. This may be described as metameric failure, where two material samples match when viewed under one light source but not another. By using some patches having the same color and shading but with different reflective and/or illumination properties, the system can distinguish between the different patches for identification. In some implementations, a given calibration chart may include keystone shaped patches that reverse the perspective of the patches. As such, any patches that are oriented away from the camera do not look too compressed and still provide a large enough area to sample.

In some implementations, a given calibration chart may be designed to take into consideration the combination of the response of a camera and the emission of the light on set. For example, if the lights are half as bright, and the camera may be twice as sensitive resulting in the same picture as if the lights were normal brightness. As such, it may be difficult to ascertain how bright the lights are.

In some implementations, a calibration chart having regular reflective patches may be used to inform the system of the relationship between the light and the camera. The backlit patches may have automatic exposure such that they get brighter under bright ambient conditions. In some implementations, an indicator on the calibration chart shows how bright they are in such a way that can be extracted from the image. As such, the system may use a single exposure to interpret the camera settings, sensitivity, color, and brightness of the light.

In various implementations, system 102 outputs all relevant information associated with calibration chart 110. For example, system 102 may output the read image parameters (e.g., RGB color values, brightness, hues, saturation, white balance, etc.) and corresponding index number for each identified patch of calibration chart 110.

In various implementations, the correction information may be used to calibrate particular camera 104, adjusting its image parameter settings for future shoots. For example, if a series of images is overexposed, patches 120 will have the same brightness. Camera 104 and possibly other cameras may be calibrated with appropriate image parameter settings. The particular adjustments or corrections may depend on the ambient lighting conditions (e.g., bright light, dark light, etc.). In some implementations, a source image that includes calibration chart 110 may be replaced with a diagnostic image for further troubleshooting and refinement of the image parameters. In some implementations, system 102 may utilize images having the same calibration chart 110, yet from multiple different cameras. This may be useful to align the image parameters associated with the different cameras so that their produced images look the same.

In various implementations, system 102 may handle various failure modes. For example, it may be difficult to identify calibration chart 110 in an image that also includes a chain link fence. In various implementations, system 110

may look at relative brightness or colors to reject. The system may gather all of the possible candidate positions of the calibration chart and then filter the positions by looking for the patches with the greatest brightness/color variation that would be consistent with calibration chart **110** being searched for. In some implementations, system **102** may crop the image **106** down in order to reduce any potential search error. System **102** may use any information derived from the read patches to determine the correction multiplier (e.g., RGB multiplier). System **102** then adjusts the RGB values based on the correction multiplier and then determines how far off patches **120** are once they have been corrected. In another failure scenario, a given patch **120** might have the wrong brightness, which may inform system **102** that a shot has been shot badly, e.g., overexposed, etc.

In various implementations, system **102** may perform the following method to facilitate an image adjustment. System **102** receives an original image from the camera, where the original image includes at least a portion of a calibration chart. System **102** derives a working image from the original image. In some implementations, the original image becomes a working image after the resolution of the original image is adjusted. For example, in some implementations, system **102** may decrease the resolution. In some implementations, system **102** may increase the resolution.

In various implementations, system **102** determines regions in the working image, where each region comprises a group of pixels having values within and/or that meet a predetermined criterion. In some implementations, determining regions based on shape includes checking an aspect ratio of a region. In some implementations, determining regions based on shape includes checking a size of a region. In some implementations, before system **102** identifies at least one region, system **102** may determine that the candidate calibration chart is at least a predetermined number of pixels in size. Otherwise, in some implementations, system **102** may crop the original image to a smaller area that includes the candidate calibration chart. System **102** may then return to the act of adjusting the resolution of the image while using the cropped original image as the original image.

In various implementations, system **102** matches at least a portion of the candidate calibration chart in the working image to a predetermined template. In various implementations, the system adjusts an orientation of the candidate calibration chart prior to the matching. For example, system **102** may determine whether the candidate calibration chart in the image is right side up, upside down, sideways, etc. based on the known, predetermined calibration chart. For example, the system may compare the colors of the identified patches in the candidate calibration chart to the colors of the known, predetermined calibration chart.

In some implementations, the matching is performed with at least a portion of the predetermined template. In some implementations, regions are not matched if a spacing of the regions is above a predetermined maximum relative distance. In some implementations, regions are not matched if a spacing of the regions is below a predetermined minimum relative distance. In some implementations, regions are not matched if a density of the regions is above a predetermined maximum density. In some implementations, regions are not matched if a density of the regions is below a predetermined maximum density.

In various implementations, system **102** may determine a confidence score (e.g., 90%, 97%, etc.) that a match is found. In some implementations, system **102** may iterate to improve the confidence score until a sufficient match is

found. System **102** may alert a predetermined person if no sufficient match is found (e.g., a match with a confidence score that is above a predetermined threshold). In some implementations, system **102** may determine a failure to match a fit of the calibration chart in the working image to the chart template. System **102** displays at least a portion of the working image on a display device. System **102** accepts a signal from a user input device to select a subset of the working image. In some implementations, system **102** may repeat one or more of the acts described herein using the selected subset of the working image.

In various implementations, system **102** filters at least a portion of the determined regions based on shape. In some implementations, the filtering at least a portion of the regions based on shape includes checking for rectangularity of a region. In various implementations, system **102** may filter portions of regions based on various information including shape (e.g., rectangularity, etc.), a number of pixels in a predetermined color range, aspect ratios of patches, etc. System **102** uses two or more of the regions to identify a candidate calibration chart in the working image. System **102** identifies at least one region within the candidate calibration chart as a patch.

In various implementations, system **102** predicts the location of one or more additional patches based on at least the identified patch. In some implementations, predicting includes using geometric extrapolation to predict patch locations from one or more identified patch locations. In some implementations, system **102** uses geometric interpolation to predict patch locations from one or more identified patch locations.

In various implementations, system **102** computes image data based on at least one pixel value from the identified patch. In various implementations, the computed image data may include image adjustment parameters. For example, in various implementations, system **102** uses at least one pixel value from the predicted patch to provide image adjustment parameters. The particular types of image data may vary, depending on the particular implementations. For example, in some implementations, the image data may include camera sensor profile.

In various implementations, system **102** computes image data based on one or more pixel values from the identified patch. In various implementations, system **102** may obtain the one or more pixel value from a middle area of the identified patch in the calibration chart to compute the image data, including one or more image parameters. System **102** may obtain pixel values from the middle area of identified patches in order to read as many pixels of the same color as possible. Pixels from the middle area of a given patch are preferred over pixels at the edge of the given patch, because pixels at the edge may have shadows, for example.

In some implementations, the image parameter may include a white balance parameter. In some implementations, the image parameter may include an exposure parameter.

Although the steps, operations, or computations may be presented in a specific order, the order may be changed in particular implementations. Other orderings of the steps are possible, depending on the particular implementation. In some particular implementations, multiple steps shown as sequential in this specification may be performed at the same time. Also, some implementations may not have all of the steps shown and/or may have other steps instead of, or in addition to, those shown herein.

In various implementations, system **102** may perform the following method to facilitate a camera adjustment. In some

implementations, system **102** receives an original image from the camera, wherein the original image includes at least a portion of a calibration chart. System **102** adjusts a resolution of the original image to create a working image. System **102** determines regions in the working image, wherein each region comprises a group of pixels having values within and/or that meet a predetermined criterion. System **102** filters at least a portion of the determined regions based on shape. System **102** uses two or more of the regions to identify a candidate calibration chart in the working image. System **102** determines that the candidate calibration chart is at least a predetermined number of pixels in size, else cropping the original image to a smaller area that includes the candidate calibration chart, then returning to the act of “adjusting a resolution,” above, while using the cropped original image as the original image. System **102** matches at least a portion of the candidate calibration chart in the working image to a predetermined template. System **102** identifies at least one region within the candidate calibration chart as a predetermined patch. System **102** uses at least one pixel value from the predetermined patch to provide an imaging parameter for adjusting the camera.

Although the steps, operations, or computations may be presented in a specific order, the order may be changed in particular implementations. Other orderings of the steps are possible, depending on the particular implementation. In some particular implementations, multiple steps shown as sequential in this specification may be performed at the same time. Also, some implementations may not have all of the steps shown and/or may have other steps instead of, or in addition to, those shown herein.

The following are additional implementations that may be used in combination with other implementations described herein. The ‘downres, process, crop, reprocess’ technique is kind of an ‘optional extra’ rather than the core of the invention. Its main purpose is to make working it possible to work quickly at low resolution without too much danger of losing quality.

For clarification, the regions are blocks of pixels that could be color chips on the chart. For example, in a “perfect” image of the 6x4 Macbeth Chart on a black background, the system may find exactly 24 regions.

The “determining region” step might find parts of the image, which are not chips of the chart (false positives) or might fail to identify a region for one of the chips on the chart (false negatives).

The point of the “matching” step is to work out which regions in the image correspond to chips on the chart, and working out which region is which chip. This makes it robust to false positives.

Having worked out that correspondence, it is possible to deduce the position of all the chips in the chart, a matrix that maps between a 2D coordinate on the surface of the chart, and a 2D pixel position.

That relationship holds well enough to predict the center of the chips, as long as the chart has not been bent, there is not too much lens distortion, and at least four regions have been identified.

This makes the matching robust to false negatives. Even though a region was not identified over a given chip, its position may be inferred from other regions that were identified.

In some implementations, there may also be a “validation” step. For example, once the values have been extracted, and values found to balance the chart, the values are checked. The error measure between what the corrected values of each patch are versus the expected value also tests

for “non-monotonic” values. In some implementations, there may be six grey patches of increasing brightness from nearly black to nearly white. If the image doesn’t also have that same progression from low to high values, then either the chart was not found correctly. Otherwise, there would be something wrong with the picture that makes it unreliable. This validation is separate from the “failure to locate chart.” System **102** may reliably find 24 regions that perfectly fit the shape of the chart, but the pixel values do not correspond to those that the chart should have. It might be the wrong sort of chart, or there could be a shadow on one of the patches.

In some implementations, system **102** may use values that the camera captured for each of the chips and known theoretical values to perform white balance, exposure correction, determination of dynamic range, linearity of sensor (gamma determination), match one camera to another, etc. In various implementations, system **102** locates a chart automatically under fairly robust conditions, reliably report when something has gone wrong.

Implementations described herein provide various benefits. For example, implementations automatically, without user intervention, correct or calibrate one or more image parameters (e.g., brightness, exposure, hue, etc.).

FIG. **8** is a block diagram of an exemplary computer system **800** for use with implementations described herein. Computer system **800** is merely illustrative and not intended to limit the scope of the claims. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. For example, computer system **800** may be implemented in a distributed client-server configuration having one or more client devices in communication with one or more server systems.

In one exemplary implementation, computer system **800** includes a display device such as a monitor **810**, computer **820**, a data entry interface **830** such as a keyboard, touch device, and the like, a user input device **840**, a network communication interface **850**, and the like. User input device **840** is typically embodied as a computer mouse, a trackball, a track pad, wireless remote, tablet, touch screen, and the like. Moreover, user input device **840** typically allows a user to select and operate objects, icons, text, characters, and the like that appear, for example, on the monitor **810**.

Network interface **850** typically includes an Ethernet card, a modem (telephone, satellite, cable, ISDN), (asynchronous) digital subscriber line (DSL) unit, and the like. Further, network interface **850** may be physically integrated on the motherboard of computer **820**, may be a software program, such as soft DSL, or the like.

Computer system **800** may also include software that enables communications over communication network **852** such as the HTTP, TCP/IP, RTP/RTSP, protocols, wireless application protocol (WAP), IEEE 902.11 protocols, and the like. In addition to and/or alternatively, other communications software and transfer protocols may also be used, for example IPX, UDP or the like. Communication network **852** may include a local area network, a wide area network, a wireless network, an Intranet, the Internet, a private network, a public network, a switched network, or any other suitable communication network, such as for example Cloud networks. Communication network **852** may include many interconnected computer systems and any suitable communication links such as hardwire links, optical links, satellite or other wireless communications links such as BLUETOOTH, WIFI, wave propagation links, or any other suitable mechanisms for communication of information. For example, communication network **852** may communicate to

one or more mobile wireless devices **856A-N**, such as mobile phones, tablets, and the like, via a base station such as wireless transceiver **854**.

Computer **820** typically includes familiar computer components such as a processor **860**, and memory storage devices, such as a memory **870**, e.g., random access memory (RAM), storage media **880**, and system bus **890** interconnecting the above components. In one implementation, computer **820** is a PC compatible computer having multiple microprocessors, graphics processing units (GPU), and the like. While a computer is shown, it will be readily apparent to one of ordinary skill in the art that many other hardware and software configurations are suitable for use with the present invention. Memory **870** and Storage media **880** are examples of tangible non-transitory computer-readable media for storage of data (or computer-readable storage media), audio/video files, computer programs, and the like. Other types of tangible media include disk drives, solid-state drives, floppy disks, optical storage media and bar codes, semiconductor memories such as flash drives, flash memories, random-access or read-only types of memories, battery-backed volatile memories, networked storage devices, Cloud storage, and the like.

FIG. **9** is a block diagram of an example visual content generation system **900**, which may be used to generate imagery in the form of still images and/or video sequences of images, according to some implementations. The visual content generation system **900** might generate imagery of live action scenes, computer generated scenes, or a combination thereof. In a practical system, users are provided with tools that allow them to specify, at high levels and low levels where necessary, what is to go into that imagery. For example, a user might be an animation artist and might use the visual content generation system **900** to capture interaction between two human actors performing live on a sound stage and replace one of the human actors with a computer-generated anthropomorphic non-human being that behaves in ways that mimic the replaced human actor's movements and mannerisms, and then add in a third computer-generated character and background scene elements that are computer-generated, all in order to tell a desired story or generate desired imagery.

Still images that are output by the visual content generation system **900** might be represented in computer memory as pixel arrays, such as a two-dimensional array of pixel color values, each associated with a pixel having a position in a two-dimensional image array. Pixel color values might be represented by three or more (or fewer) color values per pixel, such as a red value, a green value, and a blue value (e.g., in RGB format). Dimensions of such a two-dimensional array of pixel color values might correspond to a preferred and/or standard display scheme, such as 1920 pixel columns by 1280 pixel rows. Images might or might not be stored in a compressed format, but either way, a desired image may be represented as a two-dimensional array of pixel color values. In another variation, images may be represented by a pair of stereo images for three-dimensional presentations and in other variations. Some or all of an image output may represent three-dimensional imagery instead of just two-dimensional views.

A stored video sequence might include a plurality of images such as the still images described above, but where each image of the plurality of images has a place in a timing sequence, and the stored video sequence is arranged so that when each image is displayed in order, at a time indicated by the timing sequence, the display presents what appears to be moving and/or changing imagery. In one representation,

each image of the plurality of images is a video frame having a specified frame number that corresponds to an amount of time that would elapse from when a video sequence begins playing until that specified frame is displayed. A frame rate might be used to describe how many frames of the stored video sequence are displayed per unit time. Example video sequences might include 24 frames per second (24 FPS), 50 FPS, 80 FPS, or other frame rates. In some implementations, frames are interlaced or otherwise presented for display, but for the purpose of clarity of description, in some examples, it is assumed that a video frame has one specified display time and it should be understood that other variations are possible.

One method of creating a video sequence is to simply use a video camera to record a live action scene, i.e., events that physically occur and can be recorded by a video camera. The events being recorded can be events to be interpreted as viewed (such as seeing two human actors talk to each other) and/or can include events to be interpreted differently due to clever camera operations (such as moving actors about a stage to make one appear larger than the other despite the actors actually being of similar build, or using miniature objects with other miniature objects so as to be interpreted as a scene containing life-sized objects).

Creating video sequences for story-telling or other purposes often calls for scenes that cannot be created with live actors, such as a talking tree, an anthropomorphic object, space battles, and the like. Such video sequences might be generated computationally rather than capturing light from live scenes. In some instances, an entirety of a video sequence might be generated computationally, as in the case of a computer-animated feature film. In some video sequences, it is desirable to have some computer-generated imagery and some live action, perhaps with some careful merging of the two.

While computer-generated imagery might be creatable by manually specifying each color value for each pixel in each frame, this is likely too tedious to be practical. As a result, a creator uses various tools to specify the imagery at a higher level. As an example, an artist might specify the positions in a scene space, such as a three-dimensional coordinate system, might specify positions of objects and/or lighting, as well as a camera viewpoint, and a camera view plane. Taking all of those as inputs, a rendering engine may compute each of the pixel values in each of the frames. In another example, an artist specifies position and movement of an articulated object having some specified texture rather than specifying the color of each pixel representing that articulated object in each frame.

In a specific example, a rendering engine performs ray tracing wherein a pixel color value is determined by computing which objects lie along a ray traced in the scene space from the camera viewpoint through a point or portion of the camera view plane that corresponds to that pixel. For example, a camera view plane might be represented as a rectangle having a position in the scene space that is divided into a grid corresponding to the pixels of the ultimate image to be generated. In the example, a ray defined by the camera viewpoint in the scene space and a given pixel in that grid first intersects a solid, opaque, blue object, that a given pixel is assigned the color blue. Of course, for modern computer-generated imagery, determining pixel colors, and thereby generating imagery, can be more complicated, as there are lighting issues, reflections, interpolations, and other considerations.

In various implementations, a live action capture system **902** captures a live scene that plays out on a stage **904**. The

live action capture system **902** is described herein in greater detail, but might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown.

In a specific live action capture system, cameras **906(1)** and **906(2)** capture the scene, while in some systems, there might be other sensor(s) **908** that capture information from the live scene (e.g., infrared cameras, infrared sensors, motion capture (“mo-cap”) detectors, etc.). On the stage **904**, there might be human actors, animal actors, inanimate objects, background objects, and possibly an object such as a green screen **910** that is designed to be captured in a live scene recording in such a way that it is easily overlaid with computer-generated imagery. The stage **904** might also contain objects that serve as fiducials, such as fiducials **912(1)-(3)**, that might be used post-capture to determine where an object was during capture. A live action scene may be illuminated by one or more lights, such as an overhead light **914**.

During or following the capture of a live action scene, the live action capture system **902** might output live action footage to a live action footage storage **920**. A live action processing system **922** might process live action footage to generate data about that live action footage and store that data into a live action metadata storage **924**. The live action processing system **922** might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. The live action processing system **922** might process live action footage to determine boundaries of objects in a frame or multiple frames, determine locations of objects in a live action scene, where a camera was relative to some action, distances between moving objects and fiducials, etc. Where elements are detected by sensor or other means, the metadata might include location, color, and intensity of the overhead light **914**, as that might be useful in post-processing to match computer-generated lighting on objects that are computer-generated and overlaid on the live action footage. The live action processing system **922** might operate autonomously, perhaps based on predetermined program instructions, to generate and output the live action metadata upon receiving and inputting the live action footage. The live action footage can be camera-captured data as well as data from other sensors.

An animation creation system **930** is another part of the visual content generation system **900**. The animation creation system **930** might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. The animation creation system **930** may be used by animation artists, managers, and others to specify details, perhaps programmatically and/or interactively, of imagery to be generated. From user input and data from a database or other data source, indicated as a data store **932**, the animation creation system **930** might generate and output data representing objects (e.g., a horse, a human, a ball, a teapot, a cloud, a light source, a texture, etc.) to an object storage **934**, generate and output data representing a scene into a scene description storage **936**, and/or generate and output data representing animation sequences to an animation sequence storage **938**.

Scene data might indicate locations of objects and other visual elements, values of their parameters, lighting, camera location, camera view plane, and other details that a rendering engine **950** might use to render CGI imagery. For example, scene data might include the locations of several articulated characters, background objects, lighting, etc. specified in a two-dimensional space, three-dimensional space, or other dimensional space (such as a 2.5-dimensional space, three-quarter dimensions, pseudo-3D spaces, etc.) along with locations of a camera viewpoint and view plane from which to render imagery. For example, scene data might indicate that there is to be a red, fuzzy, talking dog in the right half of a video and a stationary tree in the left half of the video, all illuminated by a bright point light source that is above and behind the camera viewpoint. In some cases, the camera viewpoint is not explicit, but can be determined from a viewing frustum. In the case of imagery that is to be rendered to a rectangular view, the frustum would be a truncated pyramid. Other shapes for a rendered view are possible and the camera view plane could be different for different shapes.

The animation creation system **930** might be interactive, allowing a user to read in animation sequences, scene descriptions, object details, etc. and edit those, possibly returning them to storage to update or replace existing data. As an example, an operator might read in objects from object storage into a baking processor that would transform those objects into simpler forms and return those to the object storage **934** as new or different objects. For example, an operator might read in an object that has dozens of specified parameters (movable joints, color options, textures, etc.), select some values for those parameters and then save a baked object that is a simplified object with now fixed values for those parameters.

Rather than have to specify each detail of a scene, data from the data store **932** might be used to drive object presentation. For example, if an artist is creating an animation of a spaceship passing over the surface of the Earth, instead of manually drawing or specifying a coastline, the artist might specify that the animation creation system **930** is to read data from the data store **932** in a file containing coordinates of Earth coastlines and generate background elements of a scene using that coastline data.

Animation sequence data might be in the form of time series of data for control points of an object that has attributes that are controllable. For example, an object might be a humanoid character with limbs and joints that are movable in manners similar to typical human movements. An artist can specify an animation sequence at a high level, such as “the left hand moves from location (X1, Y1, Z1) to (X2, Y2, Z2) over time T1 to T2”, at a lower level (e.g., “move the elbow joint 2.5 degrees per frame”) or even at a very high level (e.g., “character A should move, consistent with the laws of physics that are given for this scene, from point P1 to point P2 along a specified path”).

Animation sequences in an animated scene might be specified by what happens in a live action scene. An animation driver generator **944** might read in live action metadata, such as data representing movements and positions of body parts of a live actor during a live action scene, and generate corresponding animation parameters to be stored in the animation sequence storage **938** for use in animating a CGI object. This can be useful where a live action scene of a human actor is captured while wearing mo-cap fiducials (e.g., high-contrast markers outside actor clothing, high-visibility paint on actor skin, face, etc.) and the movement of those fiducials is determined by the live



action processing system **922**. The animation driver generator **944** might convert that movement data into specifications of how joints of an articulated CGI character are to move over time.

A rendering engine **950** can read in animation sequences, scene descriptions, and object details, as well as rendering engine control inputs, such as a resolution selection and a set of rendering parameters. Resolution selection might be useful for an operator to control a trade-off between speed of rendering and clarity of detail, as speed might be more important than clarity for a movie maker to test a particular interaction or direction, while clarity might be more important than speed for a movie maker to generate data that will be used for final prints of feature films to be distributed. The rendering engine **950** might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown.

The visual content generation system **900** can also include a merging system **960** (labeled “Live+CGI Merging System”) that merges live footage with animated content. The live footage might be obtained and input by reading from the live action footage storage **920** to obtain live action footage, by reading from the live action metadata storage **924** to obtain details such as presumed segmentation in captured images segmenting objects in a live action scene from their background (perhaps aided by the fact that the green screen **910** was part of the live action scene), and by obtaining CGI imagery from the rendering engine **950**.

A merging system **960** might also read data from rule sets for merging/combining storage **962**. A very simple example of a rule in a rule set might be “obtain a full image including a two-dimensional pixel array from live footage, obtain a full image including a two-dimensional pixel array from the rendering engine **950**, and output an image where each pixel is a corresponding pixel from the rendering engine **950** when the corresponding pixel in the live footage is a specific color of green, otherwise output a pixel value from the corresponding pixel in the live footage.”

The merging system **960** might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. The merging system **960** might operate autonomously, following programming instructions, or might have a user interface or programmatic interface over which an operator can control a merging process. In some implementations, an operator can specify parameter values to use in a merging process and/or might specify specific tweaks to be made to an output of the merging system **960**, such as modifying boundaries of segmented objects, inserting blurs to smooth out imperfections, or adding other effects. Based on its inputs, the merging system **960** can output an image to be stored in a static image storage **970** and/or a sequence of images in the form of video to be stored in an animated/combined video storage **972**.

Thus, as described, the visual content generation system **900** can be used to generate video that combines live action with computer-generated animation using various components and tools, some of which are described in more detail herein. While the visual content generation system **900** might be useful for such combinations, with suitable settings, it can be used for outputting entirely live action footage or entirely CGI sequences. The code may also be

provided and/or carried by a transitory computer readable medium, e.g., a transmission medium such as in the form of a signal transmitted over a network.

According to one implementation, the techniques described herein are implemented by one or more generalized computing systems programmed to perform the techniques pursuant to program instructions in firmware, memory, other storage, or a combination. Special-purpose computing devices may be used, such as desktop computer systems, portable computer systems, handheld devices, networking devices or any other device that incorporates hard-wired and/or program logic to implement the techniques.

FIG. **10** is a block diagram of an example computer system **1000**, which may be used for implementations described herein. The computer system **1000** includes a bus **1002** or other communication mechanism for communicating information, and a processor **1004** coupled with the bus **1002** for processing information. The processor **1004** may be a general-purpose microprocessor, for example.

The computer system **1000** also includes a main memory **1006**, such as a random access memory (RAM) or other dynamic storage device, coupled to the bus **1002** for storing information and instructions to be executed by the processor **1004**. The main memory **1006** may also be used for storing temporary variables or other intermediate information during execution of instructions to be executed by the processor **1004**. Such instructions, when stored in non-transitory storage media accessible to the processor **1004**, render the computer system **1000** into a special-purpose machine that is customized to perform the operations specified in the instructions.

The computer system **1000** further includes a read only memory (ROM) **1008** or other static storage device coupled to the bus **1002** for storing static information and instructions for the processor **1004**. A storage device **1010**, such as a magnetic disk or optical disk, is provided and coupled to the bus **1002** for storing information and instructions.

The computer system **1000** may be coupled via the bus **1002** to a display **1012**, such as a computer monitor, for displaying information to a computer user. An input device **1014**, including alphanumeric and other keys, is coupled to the bus **1002** for communicating information and command selections to the processor **1004**. Another type of user input device is a cursor control **1016**, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to the processor **1004** and for controlling cursor movement on the display **1012**. This input device **1014** typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the input device **1014** to specify positions in a plane.

The computer system **1000** may implement the techniques described herein using customized hard-wired logic, one or more ASICs or FPGAs, firmware, and/or program logic, which, in combination with the computer system, causes or programs the computer system **1000** to be a special-purpose machine. According to one implementation, the techniques herein are performed by the computer system **1000** in response to the processor **1004** executing one or more sequences of one or more instructions contained in the main memory **1006**. Such instructions may be read into the main memory **1006** from another storage medium, such as the storage device **1010**. Execution of the sequences of instructions contained in the main memory **1006** causes the processor **1004** to perform the process steps described herein. In alternative implementations, hard-wired circuitry may be used in place of or in combination with software instructions.

The term “storage media” as used herein refers to any non-transitory media that store data and/or instructions that cause a machine to operate in a specific fashion. Such storage media may include non-volatile media and/or volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as the storage device **1010**. Volatile media includes dynamic memory, such as the main memory **1006**. Common forms of storage media include, for example, a floppy disk, a flexible disk, hard disk, solid state drive, magnetic tape, or any other magnetic data storage medium, a CD-ROM, any other optical data storage medium, any physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EPROM, NVRAM, any other memory chip or cartridge.

Storage media is distinct from but may be used in conjunction with transmission media. Transmission media participates in transferring information between storage media. For example, transmission media includes coaxial cables, copper wire, and fiber optics, including the wires that include the bus **1002**. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

Various forms of media may be involved in carrying one or more sequences of one or more instructions to the processor **1004** for execution. For example, the instructions may initially be carried on a magnetic disk or solid-state drive of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a network connection. A modem or network interface local to the computer system **1000** can receive the data. The bus **1002** carries the data to the main memory **1006**, from which the processor **1004** retrieves and executes the instructions. The instructions received by the main memory **1006** may optionally be stored on the storage device **1010** either before or after execution by the processor **1004**.

The computer system **1000** also includes a communication interface **1018** coupled to the bus **1002**. The communication interface **1018** provides a two-way data communication coupling to a network link **1020** that is connected to a local network **1022**. For example, the communication interface **1018** may be an integrated services digital network (“ISDN”) card, cable modem, satellite modem, or a modem to provide a data communication connection to a corresponding type of telephone line. Wireless links may also be implemented. In any such implementation, the communication interface **1018** sends and receives electrical, electromagnetic, or optical signals that carry digital data streams representing various types of information.

The network link **1020** typically provides data communication through one or more networks to other data devices. For example, the network link **1020** may provide a connection through a local network **1022** to a host computer **1024** or to data equipment operated by an Internet Service Provider (“ISP”) **1026**. The ISP **1026** in turn provides data communication services through the worldwide packet data communication network now commonly referred to as the “Internet” **1028**. The local network **1022** and the Internet **1028** both use electrical, electromagnetic, or optical signals that carry digital data streams. The signals through the various networks and the signals on the network link **1020** and through the communication interface **1018**, which carry the digital data to and from the computer system **1000**, are example forms of transmission media.

The computer system **1000** can send messages and receive data, including program code, through the network(s), the network link **1020**, and the communication

interface **1018**. In the Internet example, a server **1030** might transmit a requested code for an application program through the Internet **1028**, the ISP **1026**, the local network **1022**, and the communication interface **1018**. The received code may be executed by the processor **1004** as it is received, and/or stored in the storage device **1010**, or other non-volatile storage for later execution.

Operations of processes described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. Processes described herein may be performed under the control of one or more computer systems (e.g., the computer system **1000**) configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program including a plurality of instructions executable by one or more processors. The computer-readable storage medium may be non-transitory. The code may also be carried by any computer-readable carrier medium, such as a transient medium or signal, e.g., a signal transmitted over a communications network.

Although the description has been described with respect to particular implementations thereof, these particular implementations are merely illustrative, and not restrictive. Controls can be provided to allow modifying various parameters of the compositing at the time of performing the recordings. For example, the resolution, number of frames, accuracy of depth position may all be subject to human operator changes or selection.

Any suitable programming language can be used to implement the routines of particular implementations including C, C++, Java, assembly language, etc. Different programming techniques can be employed such as procedural or object oriented. The routines can execute on a single processing device or multiple processors. Although the steps, operations, or computations may be presented in a specific order, this order may be changed in different particular implementations. In some particular implementations, multiple steps shown as sequential in this specification can be performed at the same time.

Particular implementations may be implemented in a computer-readable storage medium for use by or in connection with the instruction execution system, apparatus, system, or device. Particular implementations can be implemented in the form of control logic in software or hardware or a combination of both. The control logic, when executed by one or more processors, may be operable to perform that which is described in particular implementations.

Some implementations are implemented as a non-transitory processor-readable medium including instructions executable by one or more digital processors. The processor-readable medium comprising one or more processor-readable instructions executable by the one or more digital processors for implementing implementations described herein.

Some implementations are implemented as processor implementable code or computer-readable code provided on a computer-readable medium. The computer-readable medium may comprise a non-transient storage medium, such as solid-state memory, a magnetic disk, optical disk, etc., or a transient medium such as a signal transmitted over a computer network.

Particular implementations may be implemented by using a programmed general purpose digital computer, by using

application specific integrated circuits, programmable logic devices, field programmable gate arrays, optical, chemical, biological, quantum or nanoengineered systems, components and mechanisms may be used. In general, the functions of particular implementations can be achieved by any means as is known in the art. Distributed, networked systems, components, and/or circuits can be used. Communication, or transfer, of data may be wired, wireless, or by any other means.

It will also be appreciated that one or more of the elements depicted in the drawings/figures can also be implemented in a more separated or integrated manner, or even removed or rendered as inoperable in certain cases, as is useful in accordance with a particular application. It is also within the spirit and scope to implement a program or code that can be stored in a machine-readable medium to permit a computer to perform any of the methods described above.

As used in the description herein and throughout the claims that follow, “a”, “an”, and “the” includes plural references unless the context clearly dictates otherwise. Also, as used in the description herein and throughout the claims that follow, the meaning of “in” includes “in” and “on” unless the context clearly dictates otherwise.

Thus, while particular implementations have been described herein, latitudes of modification, various changes, and substitutions are intended in the foregoing disclosures, and it will be appreciated that in some instances some features of particular implementations will be employed without a corresponding use of other features without departing from the scope and spirit as set forth. Therefore, many modifications may be made to adapt a particular situation or material to the essential scope and spirit.

I claim:

**1.** A computer-implemented method performed by one or more

digital processors for identifying image calibration data, the method comprising:

receiving an original image from a camera, wherein the original image includes at least a portion of a calibration chart;

deriving a working image from the original image; determining regions in the working image, wherein the regions comprise groups of pixels having values within a predetermined criterion;

analyzing two or more of the regions to identify a candidate calibration chart in the working image;

identifying at least one region within the candidate calibration chart as a patch, wherein the identifying of the at least one region is based on a color of the patch; and

predicting a location of one or more additional patches based on at least the identified patch.

**2.** The method of claim **1**, further comprising measuring a dynamic range of the camera based on one or more of a reflection and a illumination of the patch.

**3.** The method of claim **1**, wherein the determining of regions includes checking a size of a region.

**4.** The method of claim **1**, further comprising adjusting an orientation of the candidate calibration chart.

**5.** The method of claim **1**, further comprising computing image data based on at least one pixel value from the identified patch.

**6.** The method of claim **1**, further comprising computing image data based on at least one pixel value from the identified patch, wherein the at least one pixel value is obtained from a middle area of the identified patch in the calibration chart to compute the image data.

**7.** The method of claim **1**, further comprising generating an image parameter, wherein the image parameter includes a white balance parameter.

**8.** The method of claim **1**, further comprising generating an image parameter, wherein the image parameter includes an exposure parameter.

**9.** A system for identifying image calibration data, the system comprising:

one or more processors; and

logic encoded in one or more non-transitory computer-readable storage media for execution by the one or more processors and when executed operable to cause the one or more processors to perform operations comprising:

receiving an original image from a camera, wherein the original image includes at least a portion of a calibration chart;

deriving a working image from the original image; determining regions in the working image, wherein the regions comprise groups of pixels having values within a predetermined criterion;

analyzing two or more of the regions to identify a candidate calibration chart in the working image;

identifying at least one region within the candidate calibration chart as a patch, wherein the identifying of the at least one region is based on a color of the patch; and

predicting a location of one or more additional patches based on at least the identified patch.

**10.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising measuring a dynamic range of the camera based on one or more of a reflection and a illumination of the patch.

**11.** The system of claim **9**, wherein the determining of regions includes checking a size of a region.

**12.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising adjusting an orientation of the candidate calibration chart.

**13.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising computing image data based on at least one pixel value from the identified patch.

**14.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising computing image data based on at least one pixel value from the identified patch, wherein the at least one pixel value is obtained from a middle area of the identified patch in the calibration chart to compute the image data.

**15.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising generating an image parameter, and wherein the image parameter includes a white balance parameter.

**16.** The system of claim **9**, wherein the logic when executed is further operable to cause the one or more processors to perform operations comprising generating an image parameter, and wherein the image parameter includes an exposure parameter.

**17.** A non-transitory computer-readable storage medium with program instructions stored thereon, the program instructions when executed by one or more processors are operable to cause the one or more processors to perform operations comprising:

receiving an original image from a camera, wherein the original image includes at least a portion of a calibration chart;  
 deriving a working image from the original image;  
 determining regions in the working image, wherein 5  
 regions comprise groups of pixels having values within a predetermined criterion;  
 analyzing two or more of the regions to identify a candidate calibration chart in the working image;  
 identifying at least one region within the candidate cali- 10  
 bration chart as a patch, wherein the identifying of the at least one region is based on a color of the patch; and  
 predicting a location of one or more additional patches based on at least the identified patch.

**18.** The computer-readable storage medium of claim **17**, 15  
 wherein the instructions when executed are further operable to cause the one or more processors to perform operations comprising measuring a dynamic range of the camera based on one or more of a reflection and a illumination of the patch. 20

**19.** The computer-readable storage medium of claim **17**, wherein the determining of regions includes checking a size of a region.

**20.** The computer-readable storage medium of claim **17**, wherein the instructions when executed are further operable 25  
 to cause the one or more processors to perform operations comprising adjusting an orientation of the candidate calibration chart.

\* \* \* \* \*